

Shock Capturing with PDE-Based Artificial Viscosity for an Adaptive, Higher-Order Discontinuous Galerkin Finite Element Method

by

Garrett Ehud Barter

M.S., Massachusetts Institute of Technology (2004)

B.S., Massachusetts Institute of Technology (2002)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2008

© Massachusetts Institute of Technology 2008. All rights reserved.


Author
Department of Aeronautics and Astronautics
May 23, 2008

Certified by
David L. Darmofal
Associate Professor of Aeronautics and Astronautics
Thesis Supervisor

Certified by
Jaime Peraire
Professor of Aeronautics and Astronautics

Certified by
Mark Drela
Professor of Aeronautics and Astronautics

Accepted by
David L. Darmofal
Associate Department Head
Chair, Committee on Graduate Students

The seal of the Massachusetts Institute of Technology is a large, faint watermark in the background. It is a circular emblem with "MASSACHUSETTS" at the top and "INSTITUTE OF TECHNOLOGY" at the bottom. The center features a shield with a figure holding a staff, a book, and a scroll, with the year "1861" below. A banner at the bottom reads "MENS ET MANUS".

| Report Documentation Page | | | | Form Approved OMB No. 0704-0188 | |
|--|------------------------------------|-------------------------------------|---|---|---------------------------------|
| Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. | | | | | |
| 1. REPORT DATE JUN 2008 | | 2. REPORT TYPE | | 3. DATES COVERED 00-00-2008 to 00-00-2008 | |
| 4. TITLE AND SUBTITLE Shock Capturing with PDE-Based Artificial Viscosity for an Adaptive, Higher-Order Discontinuous Galerkin Finite Element Method | | | | 5a. CONTRACT NUMBER | |
| | | | | 5b. GRANT NUMBER | |
| | | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) | | | | 5d. PROJECT NUMBER | |
| | | | | 5e. TASK NUMBER | |
| | | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, Cambridge, MA, 02139 | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | | | 10. SPONSOR/MONITOR'S ACRONYM(S) | |
| | | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited | | | | | |
| 13. SUPPLEMENTARY NOTES | | | | | |
| 14. ABSTRACT <p>The accurate simulation of supersonic and hypersonic flows is well suited to high-order ($p > 1$), adaptive computational fluid dynamics (CFD). Since these cases involve flow velocities greater than the speed of sound, an appropriate shock capturing for higher-order, adaptive methods is necessary. Artificial viscosity can be combined with a higher-order discontinuous Galerkin finite element discretization to resolve a shock layer within a single cell. However, when a nonsmooth artificial viscosity model is employed with an otherwise higher-order approximation, element-to-element variations induce oscillations in state gradients and pollute the downstream flow. To alleviate these difficulties, this work proposes a new, higher-order, state-based artificial viscosity with an associated governing partial differential equation (PDE). In the governing PDE, the shock sensor acts as a forcing term, driving the artificial viscosity to a non-zero value where it is necessary. The decay rate of the higher-order solution modes and edge-based jumps are both shown to be reliable shock indicators. This new approach leads to a smooth, higher-order representation of the artificial viscosity that evolves in time with the solution. For applications involving the Navier-Stokes equations, an artificial dissipation operator that preserves total enthalpy is introduced. The combination of higher-order, PDE-based artificial viscosity and enthalpy-preserving dissipation operator is shown to overcome the disadvantages of the non-smooth artificial viscosity. The PDE-based artificial viscosity can be used in conjunction with an automated grid adaptation framework that minimizes the error of an output functional. Higher-order solutions are shown to reach strict engineering tolerances with fewer degrees of freedom. The benefit in computational efficiency for higher-order solutions is less dramatic in the vicinity of the shock where errors scale with $O(h/p)$. This includes the near-field pressure signals necessary for sonic boom prediction. When applied to heat transfer prediction on unstructured meshes in hypersonic flows, the PDE-based artificial viscosity is less susceptible to errors introduced by poor shock-grid alignment. Surface heating can also drive the output-based grid adaptation framework to arrive at the same heat transfer distribution as a well-designed structured mesh.</p> | | | | | |
| 15. SUBJECT TERMS | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT Same as Report (SAR) | 18. NUMBER OF PAGES 143 | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT unclassified | b. ABSTRACT unclassified | c. THIS PAGE unclassified | | | |

Shock Capturing with PDE-Based Artificial Viscosity for an Adaptive, Higher-Order Discontinuous Galerkin Finite Element Method

by
Garrett Ehud Barter

Submitted to the Department of Aeronautics and Astronautics
on May 23, 2008, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The accurate simulation of supersonic and hypersonic flows is well suited to higher-order ($p > 1$), adaptive computational fluid dynamics (CFD). Since these cases involve flow velocities greater than the speed of sound, an appropriate shock capturing for higher-order, adaptive methods is necessary.

Artificial viscosity can be combined with a higher-order discontinuous Galerkin finite element discretization to resolve a shock layer within a single cell. However, when a non-smooth artificial viscosity model is employed with an otherwise higher-order approximation, element-to-element variations induce oscillations in state gradients and pollute the downstream flow. To alleviate these difficulties, this work proposes a new, higher-order, state-based artificial viscosity with an associated governing partial differential equation (PDE). In the governing PDE, the shock sensor acts as a forcing term, driving the artificial viscosity to a non-zero value where it is necessary. The decay rate of the higher-order solution modes and edge-based jumps are both shown to be reliable shock indicators. This new approach leads to a smooth, higher-order representation of the artificial viscosity that evolves in time with the solution. For applications involving the Navier-Stokes equations, an artificial dissipation operator that preserves total enthalpy is introduced. The combination of higher-order, PDE-based artificial viscosity and enthalpy-preserving dissipation operator is shown to overcome the disadvantages of the non-smooth artificial viscosity.

The PDE-based artificial viscosity can be used in conjunction with an automated grid adaptation framework that minimizes the error of an output functional. Higher-order solutions are shown to reach strict engineering tolerances with fewer degrees of freedom. The benefit in computational efficiency for higher-order solutions is less dramatic in the vicinity of the shock where errors scale with $\mathcal{O}(h/p)$. This includes the near-field pressure signals necessary for sonic boom prediction. When applied to heat transfer prediction on unstructured meshes in hypersonic flows, the PDE-based artificial viscosity is less susceptible to errors introduced by poor shock-grid alignment. Surface heating can also drive the output-based grid adaptation framework to arrive at the same heat transfer distribution as a well-designed structured mesh.

Thesis Supervisor: David L. Darmofal

Title: Associate Professor of Aeronautics and Astronautics

Acknowledgments

This thesis represents nearly four years of work and study. I would not have made it through one year, let alone four, without the professional guidance and personal comradery of others.

First and foremost, I would like to express my deepest gratitude to my advisor, Prof. David Darmofal. I have learned a great deal from his experience, intuition and thoroughness. He has guided me in bringing an abstract idea to the level of rigorous research. More importantly, he supported and believed in me during both the highs and lows of the project. The merits of this work are a credit to his mentorship while any faults are my own. In addition to my advisor, the other members of my thesis committee, Prof. Jaime Peraire and Prof. Mark Drela, also deserve considerable recognition for keen insights and criticisms of my work along the way. Their experience and ability to foresee problems were extremely valuable. I would also like to thank my thesis readers, Prof. Graham Candler and Dr. Peter Gnoffo, for their reviews. This thesis has benefited from their commentary.

This project found its direction thanks to the work of Douglas Quattrochi. His careful, patient and meticulous investigation of the hypersonic flow over a cylinder painted a clear picture of the need for a different artificial viscosity model. The contributions in this thesis would never have been realized without his research. Loretta Trevino deserves notable credit for banging on my code until it worked and laying the groundwork for the hypersonic ramp case. I also owe a great debt to my Project-X collaborators, especially Todd Oliver, Krzysztof Fidkowski and Michael Park, as well as Laslo Diosady, JM Modisette and Joshua Krakos. Sometimes the obstacles I encountered could only be overcome with their added insight and willingness to help. Similarly, I benefited from the fruits of other people's labors who wrote the vital organs of Project-X. In addition to those mentioned previously, I would like to thank Matthieu Serrano, Michael Brasher, James Lu, Paul Nicholson, Eric Liu, Eleanor Lin, Peter Whitney, Shannon Cheng, Jean-Baptiste Brachet, Huafei Sun and Masayuki Yano. The conversations and advice from Robert Haines, Per-Olof Persson, Benjamin Kirk, Thomas Richter and Prof. Karen Willcox all proved useful at various points in my research. I apologize for not being more personal in my gratitude. Finally, no research would get done in the lab without the undervalued contributions of Jean Sofronas and Ping Lee. I am grateful for their countless favors and efforts on my behalf.

I exit graduate school having gained valuable experience in both research and computer systems administration. I appreciate the opportunity Bob gave me to maintain and grow the lab resources while helping others be productive. What I have learned about systems administration is just a tiny fraction of the vast amount of knowledge passed on to me by Bob, David Smith and Victor Garzon.

One of the more rewarding activities I have been involved in during graduate school, as a way to give back to the Department that has given me so much, is the Graduate Association of Aeronautics and Astronautics (GA³). I truly enjoyed working with Jessica Marquez, Jessica Townsend and Ryan Peoples in establishing the group and setting its direction. Without the help of initial members, Kevin Duda, Geoff Huntington, Steve Paschall, Emily Craparo, Daniel Kwon, Noah Warner, Ryan Whittaker, Anotine Jerusalem, Hayley Davidson, Theresia Becker, Jessica Edmonds, Phil Ferguson, Farmey Joseph, Jon Histon, Kristen Bethke, Katie Weiss, Kelly Klima, Sajjad Shaikh Matin and David Benson, GA³ would not have survived. I am proud to have been a part of this organization.

The comradery of others made the many hours I spent in ACDL bearable. In no particular order I have enjoyed the company of Yann Deremaux, David Gratton, Matthieu Serrano, Jerome Lavain, Jean-Baptiste Brachet, Curran Crawford, Victor Garzon, Vince

Sidwell, Dan King, Ryan Peoples, Sean Bradshaw, Joe Alescio, Sudeep Lahiri, Hector Ciria, Marc Santi, Chris Hynes, Tony Lau, Vivian Sho, Kristin Jonker, Caroline Lamb, Jean Collin, Mike Brasher, Paul Nicholson, Todd Oliver, Mark Monroe, Shana Diez, Krzysztof Fidkowski, Mike Park, Shannon Cheng, Thomas Richter, Per-Olaf Persson, David Lazzara, Dan Dicara, Svein Hovland, Nuria Pares, Peter Whitney, James Lu, Luis Cueto, Yolanda Vidal-Segui, David Walfisch, Doug Quattrochi, David Willis, Tan Bui, Theresa Robinson, Leia Stirling, Doug Allaire, Omar Bashir, Tudor Masek, Emily Israeli, Alejandra Uranga, Sonja Wogrin, JM Modisette, Laslo Diosady, Josh Krakos and others.

I am thankful to those with whom I spent a great deal of time with outside of lab. My roommates through graduate school, Mike Brasher, Mike Roberts, Elie Krevat and Todd Oliver. Todd- I enjoyed our time together as roommates and watching each other obtain degrees, get married and lose our hair. There are also two people who have been with me from the first week of MIT to the last- Dan Kwon and Vijay Divi. You are terrific friends, thank you.

I of course would like to thank my family for providing me with the support, encouragement and independence in reaching this milestone. Mom, Dad, Mickey and Devra- I love you all very much. I also enjoyed having family in the Boston area who took care of me during weekends and holidays. Thank you to Peggy, Gene, Jessica, David, Joshua and Ben for welcoming me into your home. I will miss your company.

To my wife, Elke, who has earned her doctorate degree in support, encouragement, patience and many other virtues these past few years. To say that I could not have completed this project without you is an understatement. Words are an inadequate means for expressing my love and appreciation.

I would be remiss if I did not acknowledge the financial support I received during the course of this project. This work was made possible by the following sources,

- National Defense Science and Engineering Graduate Fellowship
- U. S. Air Force Research Laboratory (USAF-3306-03-SC-0001)
- The Boeing Company
- National Aeronautics and Space Administration (Co-op Agreement NNX07AC70A)

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 15 |
| 1.1 | Motivation | 15 |
| 1.1.1 | Sonic Boom Prediction | 15 |
| 1.1.2 | Hypersonics | 18 |
| 1.2 | Thesis Objective | 20 |
| 1.3 | Background | 20 |
| 1.3.1 | Higher-Order Methods | 20 |
| 1.3.2 | Shock Capturing | 21 |
| 1.3.3 | Error Estimation and Grid Adaptation | 25 |
| 1.4 | Thesis Overview | 27 |
| 2 | Discontinuous Galerkin Discretization and the Compressible Navier-Stokes Equations | 29 |
| 2.1 | Discontinuous Galerkin Finite Elements | 29 |
| 2.1.1 | Solution and Geometry Interpolation | 31 |
| 2.1.2 | Solution Method | 32 |
| 2.2 | Compressible Navier-Stokes Equations | 32 |
| 2.3 | Artificial Viscosity Matrix | 33 |
| 2.3.1 | Numerical Diffusion for Constant Total Enthalpy | 34 |
| 3 | Motivation for Smooth Artificial Viscosity | 37 |
| 3.1 | Vanishing Viscosity and Conservation Laws | 37 |
| 3.1.1 | Burgers' Equation Example | 39 |
| 3.2 | One-Dimensional Results | 41 |
| 3.3 | Adjoint Analyses | 44 |
| 3.3.1 | Design Variable Sensitivity Error | 47 |
| 3.4 | Multiple Dimension Issues | 48 |
| 4 | A PDE-Based Artificial Viscosity Model | 51 |
| 4.1 | Non-Smooth Artificial Viscosity | 51 |
| 4.2 | PDE-based Artificial Viscosity | 51 |
| 4.2.1 | Modified System of Equations | 53 |
| 4.2.2 | Green's Function Behavior | 54 |
| 4.3 | Shock Indicators | 56 |
| 4.3.1 | Resolution Indicator | 57 |
| 4.3.2 | Jump Indicator | 58 |
| 4.4 | Artificial Viscosity Model Comparisons | 59 |

| | | |
|----------|---|------------|
| 4.4.1 | Convergence Rate Accuracy | 59 |
| 4.4.2 | Transonic flow: NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$ | 62 |
| 4.4.3 | Supersonic flow: Half-cylinder, $M_\infty = 4$ | 64 |
| 5 | Output-Based Grid Adaptation with Shocks | 69 |
| 5.1 | Error Estimation | 69 |
| 5.1.1 | Error Estimation with Artificial Viscosity | 72 |
| 5.2 | Adaptation Mechanics | 74 |
| 5.3 | Two Dimensional Results | 77 |
| 5.3.1 | Supersonic Flow: Compression Ramp, $M_\infty = 12$ | 77 |
| 5.3.2 | Transonic flow: NACA 0012, $M_\infty = 0.95$, $\alpha = 0^\circ$ | 83 |
| 5.3.3 | Transonic flow: NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$ | 89 |
| 5.3.4 | Supersonic flow: NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ | 92 |
| 6 | Hypersonic Applications | 101 |
| 6.1 | Compression Ramp, $M_\infty = 11.68$, $\text{Re} = 246,636$ | 101 |
| 6.1.1 | Structured Grid Results | 103 |
| 6.1.2 | Adaptation Results | 108 |
| 6.2 | Half Cylinder, $M_\infty = 17.605$, $\text{Re} = 376,930$ | 110 |
| 6.2.1 | Previous Research | 111 |
| 6.2.2 | Discontinuous Galerkin Approach | 114 |
| 6.2.3 | Structured Grid Results | 114 |
| 6.2.4 | Unstructured Grid Results, Two Dimensions | 115 |
| 6.2.5 | Unstructured Grid Results, Three Dimensions | 118 |
| 6.2.6 | 2D Adaptation | 120 |
| 7 | Conclusions | 127 |
| 7.1 | Summary and Contributions | 127 |
| 7.2 | Future Work | 129 |
| A | Dual Consistency of Nonlinear Viscosity | 131 |
| A.1 | Dual Consistency Preliminaries | 131 |
| A.2 | Dual Consistency of the Non-Linear Poisson Equation | 132 |
| | Bibliography | 135 |

List of Figures

| | | |
|-----|---|----|
| 1-1 | Stages in computational modeling for sonic boom generation and propagation (from [25]). | 17 |
| 1-2 | Shock layer resolution in DG FEM for increasing polynomial orders. | 24 |
| 3-1 | Traveling wave solution of Burgers' equation with vanishing viscosity and shock width diagram. | 40 |
| 3-2 | Distributions of piecewise-constant and Gaussian artificial viscosity as applied to the 1D modified Burgers equation. | 41 |
| 3-3 | Comparison of piecewise-constant and Gaussian viscosity solutions for modified Burgers equation across three different viscosity amplitudes (40 elements, $p = 6$). . | 43 |
| 3-4 | Distributions of expanded piecewise-constant and Gaussian artificial viscosity as applied to the 1D modified Burgers equation. | 44 |
| 3-5 | Comparison of expanded piecewise-constant and Gaussian viscosity solutions for modified Burgers equation across three different viscosity amplitudes (40 elements, $p = 6$). | 45 |
| 3-6 | Grid convergence study of variable sensitivity errors computed via the adjoint for piecewise-constant and Gaussian distributions of viscosity. | 48 |
| 3-7 | Mach number contours and total pressure measurements along a line behind the bow shock across two grid refinements for a $p = 3$ solution of a 2D flow around a cylinder, $M_\infty = 4$ | 49 |
| 4-1 | Green's function for the 1D model of the artificial viscosity equation. | 55 |
| 4-2 | Boundary condition impact upon Green's function behavior for a source near the domain boundary. | 56 |
| 4-3 | Gaussian bump domain and mesh for smooth flow, shock indicator accuracy study (1600 elements). | 60 |
| 4-4 | Grid convergence rates of global entropy norm for inviscid flow over a Gaussian bump, $M_\infty = 0.5$, $\alpha = 0^\circ$ with non-smooth (NS) and PDE-based artificial viscosity models. | 61 |
| 4-5 | L^1 grid convergence rates for 1D modified Burgers equation of a forcing function with discontinuity with both non-smooth (NS) and PDE-based artificial viscosity. . | 63 |
| 4-6 | Comparison of non-smooth and PDE-based artificial viscosity solutions for modified Burgers equation (40 elements, $p = 6$). | 64 |
| 4-7 | Inviscid $p = 5$ solution with resolution shock indicator of a NACA 0012 airfoil, $M_\infty = 0.8$ and $\alpha = 1.25^\circ$ | 65 |
| 4-8 | Three grids across two uniform grid refinements used for inviscid flow over a 2D half cylinder. | 66 |

| | | |
|------|---|----|
| 4-9 | Comparison of viscosity models and shock capturing performance for a $p = 3$ solution of a 2D flow around a cylinder at Mach 4, resolution shock indicator (contour plots are shown for the intermediate mesh, vLH is van Leer-Hänel flux function). | 67 |
| 4-10 | Comparison of total pressure along a measurement line behind the bow shock across two grid refinements for a $p = 3$ solution of a 2D flow around a cylinder at Mach 4 (resolution shock indicator with van Leer-Hänel flux function and $\tilde{\mathbf{A}}_\epsilon$ viscosity matrix). | 68 |
| 5-1 | Comparison of exact solution approximations near a shock for the modified 1D Burgers equation, $p = 2$ solution with 50 elements. | 71 |
| 5-2 | Additional comparison of exact solution approximations near a shock and error estimate effectivities for the modified 1D Burgers equation, $p = 2$ solution with 50 elements, with discontinuous solution. | 75 |
| 5-3 | Inviscid flow over a 15° wedge, $M_\infty = 12$ | 77 |
| 5-4 | Nested structured meshes of a 15° wedge. | 77 |
| 5-5 | Entropy norm error and shock width convergence for inviscid flow over a 15° wedge, $M_\infty = 12$ | 78 |
| 5-6 | Zoom of Mach number contours at the shock overlayed with the mesh for inviscid flow over a 15° wedge, $M_\infty = 12$ ($p = 3$ solution on Grid 1). | 79 |
| 5-7 | Final adaptation grids for inviscid flow over a 15° wedge, $M_\infty = 12$ | 80 |
| 5-8 | Error and functional convergence histories for inviscid flow over a 15° wedge, $M_\infty = 12$ | 81 |
| 5-9 | Shock width convergence of nested structured meshes and adapted grids for inviscid flow over a 15° wedge, $M_\infty = 12$ | 82 |
| 5-10 | Effectivity histories for inviscid flow over a 15° wedge, $M_\infty = 12$ | 83 |
| 5-11 | Region of modified regularity due to artificial viscosity for inviscid flow over a 15° wedge, $M_\infty = 12$ | 83 |
| 5-12 | Near-field and far-field contour plots of Mach number and density adjoint for drag, inviscid transonic flow over NACA 0012, $M_\infty = 0.95$ (fine mesh, truth solution, $p = 3$). | 85 |
| 5-13 | Near-field view of initial and final adaptation meshes for inviscid transonic flow over NACA 0012, $M_\infty = 0.95$ | 86 |
| 5-14 | Functional and error histories of adaptation process for inviscid transonic flow over NACA 0012, $M_\infty = 0.95$ | 87 |
| 5-15 | Sensitivity of higher-order adaptation sequence to starting mesh resolution for inviscid transonic flow over NACA 0012, $M_\infty = 0.95$ | 88 |
| 5-16 | Mach and x -momentum adjoint for lift contours for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$ | 89 |
| 5-17 | Initial and final adaptation meshes for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$ | 90 |
| 5-18 | Functional and error histories of adaptation process for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$ | 91 |
| 5-19 | Regions of modified regularity due to artificial viscosity for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$ | 92 |
| 5-20 | Mach and density adjoint for drag contours of viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ | 93 |
| 5-21 | Initial and final meshes, for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ (adaptation for drag). | 94 |
| 5-22 | Functional and error histories of adaptation process for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ (adaptation for drag). | 95 |

| | | |
|------|---|-----|
| 5-23 | Near-field and far-field contour plots of x -momentum and its adjoint for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ (adaptation for far-field pressure). | 96 |
| 5-24 | Region of artificial viscosity greater than kinematic viscosity for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ | 97 |
| 5-25 | Initial and final meshes, for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ (adaptation for far-field pressure). | 98 |
| 5-26 | Functional and error histories of adaptation process for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ (adaptation for far-field pressure). | 99 |
| 6-1 | Geometry and boundary conditions for hypersonic compression ramp problem. . . | 101 |
| 6-2 | Contour plots of $p = 3$ solution of flow over compression ramp with $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 102 |
| 6-3 | Comparison of computational and experimental results for the compression ramp problem obtained by Lillard and Dries [84] and Kirk [81]. | 104 |
| 6-4 | Computational meshes used for grid convergence study in compression ramp problem. | 105 |
| 6-5 | Region of artificial viscosity greater than kinematic viscosity of $p = 3$ solution of flow over compression ramp with $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 105 |
| 6-6 | Mach contour plot convergence of $p = 1-3$ on Grids 1 and 2 for flow over compression ramp with $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 105 |
| 6-7 | Surface plots on a given mesh from grid convergence study of flow over a compression ramp, $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 106 |
| 6-8 | Grid converged surface plots for flow over a compression ramp, $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 107 |
| 6-9 | Mach number along line extending upwards from ramp corner from grid convergence study of flow over a compression ramp, $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 108 |
| 6-10 | Surface plots from final adapted solution of flow over a compression ramp, $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 109 |
| 6-11 | Final unstructured, adapted mesh, error history and functional history of adaptation process for flow over a compression ramp, $M_\infty = 11.68$, $\text{Re} = 246, 636$ | 110 |
| 6-12 | Computational mesh and results obtained by Gnoffo and White [55]. | 111 |
| 6-13 | Computational meshes used by Nompelis et al. [98]. | 112 |
| 6-14 | Surface heat transfer and pressure coefficient obtained by Nompelis et al. [98]. . . | 113 |
| 6-15 | Structured grids used for 2D half cylinder grid convergence study. | 114 |
| 6-16 | Surface plots on a given mesh from grid convergence study of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ | 116 |
| 6-17 | Region of artificial viscosity greater than kinematic viscosity for flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ (taken from $p = 3$ solution on finest structured mesh). | 117 |
| 6-18 | Structured-unstructured hybrid grids used for 2D half cylinder test case. | 117 |
| 6-19 | Contour plots of $p = 3$ solution of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ (taken from Hybrid Grid 0). | 118 |
| 6-20 | Cylinder surface plots of $p = 3$ solution of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ (all 5 grids shown). | 119 |
| 6-21 | Mach contours at the bow shock of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ ($p = 3$ solution). | 120 |
| 6-22 | Structured-unstructured hybrid grid used for 3D extruded half cylinder test case. . | 121 |

| | | |
|------|---|-----|
| 6-23 | Contour plots of $p = 3$ solution of flow over 3D extruded half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ | 121 |
| 6-24 | Cylinder surface plots of $p = 3$ solution of flow over 3D extruded half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ | 122 |
| 6-25 | Initial and final adapted mesh, cylinder contour plots, error history and functional convergence for adaptation of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$, $p = 2$ | 124 |
| 6-26 | Stagnation line plots of $p = 3$ solution of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ (taken from Grid 0). | 125 |
| 6-27 | Cylinder surface plots of adapted solution of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ | 125 |
| 6-28 | Cylinder surface plots of adapted solution of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376, 930$ | 126 |

List of Tables

| | | |
|-----|--|-----|
| 2.1 | Viscous fluxes | 31 |
| 3.1 | Global L^2 error norm and H^1 error norm outside of shock layer comparisons of piecewise-constant, expanded piecewise-constant, and Gaussian viscosity solutions for modified Burgers equation across three different viscosity amplitudes (40 elements, $p = 6$) | 44 |
| 5.1 | Scaling factor of global error with respect to polynomial order for shock dominated flows. | 80 |
| 6.1 | Freestream and boundary values for compression ramp test case. | 102 |
| 6.2 | Freestream and boundary values for half-cylinder test case. | 111 |

Chapter 1

Introduction

In the past decades, computational resources and algorithms have matured to a state such that numerical modeling is an essential component of engineering design and analysis. This is certainly true for computational fluid dynamics (CFD), which has grown into the ability to solve flow fields with sophisticated geometries and complex physical processes. While experimental measurements will always have a role in the design process, CFD offers advantages in terms of cost, test time, ease of use, and quality of output data. Nevertheless, despite the advances in the usage and capabilities of CFD, there is still room for improvement.

One area of CFD growth is in the development of higher-order accurate schemes and their application to an expanding diversity of flow regimes and problems. Two such flow regimes, supersonic and hypersonic flow, serve as the motivation for this work. In both instances, higher-order CFD solutions, with efficient adaptive capabilities based on a functional output, can advance the state-of-the-art in flow field modeling and predictive capabilities. Since these applications involve flow velocities greater than the speed of sound, where shock waves develop, the focus of this work is on a shock capturing methodology for higher-order and adaptive methods.

1.1 Motivation

Higher-order solutions, with efficient adaptive capabilities based on functional outputs, coupled with a robust and accurate shock capturing methodology offer advantages in many applications. One example is the accurate prediction of sonic boom footprints, extrapolated from CFD simulations of the near-field flow around an aircraft. Another example is accurate estimates of heating and shear and pressure forces on a body in hypersonic flow.

1.1.1 Sonic Boom Prediction

The sonic boom phenomenon is one of the chief factors hindering the use of supersonic flight over land and populated areas. In 2001, the National Research Council Committee on Breakthrough Technology for Commercial Supersonic Aircraft investigated the feasibility of commercial supersonic flight and made recommendations to NASA to realize that

goal [37]. They determined that sonic boom mitigation was the key technological barrier to the development of a supersonic business jet. Fortunately, recent advances in sonic boom reduction technology might enable overland supersonic cruise for both military and commercial applications [36]. The potential benefits of quiet, supersonic cruise include reduced travel time for business or cargo and rapid response or strike capability for the military [36]. Thus, strong motivation exists from both the civilian and military communities to minimize the acoustic footprint of supersonic aircraft.

The recent advances that give promise to the future of supersonic flight are both new technologies intended to mitigate sonic boom intensity and also new design capabilities. One of the recommendations made by the National Research Council was to bring high-fidelity analysis of new concepts and technologies to the early stages of conceptual design in a multi-disciplinary environment. In this way, the sonic boom signature can be an integral design metric. However, in order for these new technologies and design processes to become accepted engineering tools, the modeling of the sonic boom phenomenon must be credible. Thus, there is a specific need for adaptive CFD in the design process for sonic boom reduction that can be used in conjunction with boom propagation codes [37].

Modeling Approach

Plotkin [112] described the standard approach to sonic boom modeling as a division of the problem into three different stages, shown in Figure 1-1. The first stage is the *near-field* and comprises the disturbances created by the possibly complex aircraft geometry as it travels at supersonic speeds. In this stage, atmospheric gradients can be ignored compared to the strong disturbances caused by the body. Furthermore, due to the geometric complexities and the strength of the flow disturbances, CFD simulations are well-suited for near-field modeling. The pressure perturbations created in the near-field then propagate through the real atmosphere, where changes in acoustic impedance and non-linear effects distort the signature. This process is commonly modeled with ray tracing and geometric acoustics [110]. The variations in the pressure signature are significant enough such that high pressure peaks propagate faster than the low pressure troughs. The *mid-field* region describes the area where this non-linear distortion occurs, but where the signal still retains features of the aircraft geometry. The *far-field* refers to the region where the acoustic signature approaches an asymptotic state, typically an N-wave, which can be computed with Whitham's rule [112].

Once the far-field pressure signature is estimated, it must be converted into practical, human-weighted metrics. This process must take into account ground absorption, ground reflection, and human ear sensitivity to different parts of the frequency spectrum [110]. Additionally, outdoor annoyance can depend on different signal characteristics than indoor annoyance [123]. The N-wave structure of many boom signatures, with two strong shocks, can be particularly loud in the weighted metrics. Thus, much research has been devoted to supersonic aircraft that produce alternate wave shapes, such as *ramps* or *flattops* [25, 123].




Figure 1-1: Stages in computational modeling for sonic boom generation and propagation (from [25]).

CFD Challenges

Ideally, the near-field CFD prediction is carried out to a far enough distance from the aircraft that cross flow diffraction effects are negligible and the pressure perturbations can be modeled as radiating sources [105, 112]. However, this near-field terminus can sometimes be located many body lengths away from the aircraft, making the CFD solution computationally prohibitive. Also, the numerical dissipation of a CFD scheme can overly attenuate the pressure signature. Some researchers have developed models to translate CFD solutions close to the body to pressure signals that can be handed off to far-field propagation codes [105, 116]. Others have used grid adaptation to accurately capture the pressure signature at the near-field terminus. Choi et al. [27] found that to achieve good agreement in the far-field estimation of noise metrics compared to experimental data, their near-field adapted grids required on the order of 10^7 nodes and tetrahedra. For relatively simple shock structures, computational costs can be reduced using shock fitting techniques [80, 104].

Higher-order, adaptive CFD methods are well positioned to improve the current state-of-the-art in near-field sonic boom prediction. Higher-order methods are recognized in aeroacoustics for their ability to capture complex features across the frequency spectra in a computationally efficient manner [6]. Robust and efficient adaptive methods, with quantifiable error estimates, are a key component of future variable fidelity, multi-disciplinary, and multi-objective optimization techniques that are necessary in the design of next-generation supersonic aircraft [36]. The combination of higher-order and output-based grid adaptation promises accurate sonic boom predictive capability while at the same time being computationally efficient.

1.1.2 Hypersonics

The onset of hypersonic flow is dependent on the flow conditions and body geometry. In general, hypersonic flow is dominated by a few characteristics that emerge as important flow phenomenon in the range of Mach numbers between 3–7. These notable characteristics are: [76, 117]

Thin shock layers: For increasing Mach numbers, oblique shock angles over slender bodies become smaller and smaller. Thus, shock waves tend to hug closely to the body geometry at hypersonic speeds, creating small layers of flow between the shock and the body itself. This leads to shock wave interactions with other flow phenomenon, such as secondary shocks, shear layers, and boundary layers.

Entropy layers and gradients For flows over blunt bodies, a highly-curved bow shock ahead of the vehicle creates a non-uniform entropy field behind the shock. On a streamline close to the vehicle nose, the bow shock appears as a strong normal shock. In contrast, a streamline far away from the nose might encounter a weak oblique shock instead. Via Crocco’s Theorem, these entropy gradients behind the shock induce vorticity as well.

Viscous interaction: Hypersonic flows, in which the freestream kinetic energy dominates over the static thermal energy, are slowed to zero velocity within the boundary layer to satisfy the no-slip boundary condition. The resulting heat release markedly increases the heat transfer and skin friction values on the surface. The increased temperature in the boundary layer decreases the density and also increases the viscosity coefficient via Sutherland’s law. As a result, the displacement thickness of hypersonic boundary layers is larger than for low Mach numbers at the same Reynolds number. This causes notable changes to the effective body geometry that the external flows sees and also increases the likelihood of the boundary layer intersecting shock waves in the flow field.

High temperature effects: The kinetic energy conversion in hypersonic flows due to viscous dissipation and/or strong shock transitions leads to high gas temperatures. The temperature can increase until the thermal energy of the gas is comparable to the energies associated with molecular processes such as vibrational excitation, disassociation, and ionization. The gas, therefore, no longer behaves as thermally and calorically perfect and must instead be considered chemically reacting. Sometimes the reaction time scale is on the same order as the those in the flow field and the gas must further be considered to be non-equilibrium flow. Finally, the temperature in hypersonic flows can become so high that radiative heat transfer becomes an important contributor to the overall heat load to a body.

Today, hypersonic flight is commonly realized by rocket-powered launch vehicles ascending through the upper-atmosphere and by unpowered reentry vehicles descending through the atmosphere of earth or other celestial bodies [111]. Recent advances in air-breathing propulsion for hypersonic flows, such as SCRAMJET technology and the successful X-43

research mission, suggest promising advances in hypersonic transport [20]. However, shortcomings in the scientific community’s understanding of the fundamental physical processes involved in hypersonic flight and ability to simulate these processes are barriers to reliable, reusable engineering systems that operate in the hypersonic regime [145]. One of NASA’s aeronautics research goals is to “develop predictive capabilities enabling both the civilian and military communities to build hypersonic systems that meet their specific needs” [111].

CFD Challenges

Due to the complex physical process, the extremes of temperature, pressure, and density, and the non-linear governing principles, the use of computation to simulate hypersonic flow fields is critical for engineering applications. For instance, on any reentry vehicle, the accurate prediction of the heat transfer distribution on a body over the entire flight path is an essential ingredient of the design process. A vehicle’s external shape and thermal protection system design are directly impacted by the aerothermodynamic simulation of the reentry flow field [54]. Unfortunately, the large uncertainties resulting from poor predictive accuracy of the aerothermodynamics, structural interaction, and material properties lead to large engineering margins in the design process, limiting performance, and increasing cost [111]. One recent example of the impact of this uncertainty on risk management was the addition of an extra space-walk on shuttle flight STS-114 to remove gap filler protruding from the tiles of the shuttle thermal protection system. Low confidence in the ability to simulate the impact of the gap filler on reentry heating and boundary layer transition suggested that the added space walk was deemed to be lower risk than leaving the gap filler in place [111].

The complex physical phenomenon and the large spectrum of spatial and temporal scales in hypersonic flow make the development of numerical simulations challenging. Accurate prediction of surface heating requires identification of transition locations from laminar to turbulent flow, inclusion of thermal and chemical non-equilibrium effects, radiative heat transfer behind strong shocks in the thin shock layer, and dynamic ablation contributions. Additionally, a hypersonic flow field might include regimes of both continuum and rarefied gas, requiring the use of very different physical models. Finally, the shock/shock and shock/boundary layer interactions can result in unsteady flow behavior necessitating time accurate computations.

Given the difficulties in developing a complete numerical tool suited for hypersonic applications, it is tempting to devise simple test problems, with limited flow complexities, for hypersonic CFD validation studies. These simplified problems, such as an axisymmetric body in non-reacting, laminar flow, could be combined with experimental data to construct a series of validation problems of varying complexity. However, obtaining experimental data in ground test facilities can also be difficult for hypersonic flows. The flow visualization and measurement techniques that are robust enough to withstand a hypersonic flow environment are relatively limited [111]. Furthermore, it is difficult to ensure that the flow field in the test section is perfectly quiet and steady [51], and, in some cases, accurate prediction of heat transfer on a body in the experimental test section requires computational simulation of the flow through the entire facility [23].

Even for problems with simplified physical processes, the quality of a hypersonic flow CFD simulation still depends strongly on the quality of the computational mesh. Unstructured meshes are superior to structured meshes in their ability to conform to complex body geometries for initial mesh generation and in their ability to adapt to the many flow features present in the flow field via anisotropic adaptation [51]. Yoon et al. [145] claims that unstructured meshes offer the greatest promise for the development of a robust, computational aerothermodynamic tool. However, the solution quality using unstructured meshes for current state-of-the art codes is far inferior to that of structured meshes. The poor solution quality manifests itself even in symmetric, simplified test cases with poor prediction of peak heat transfer rates and asymmetric surface heat transfer distributions. The problem stems from the misalignment of the unstructured mesh with the strong shocks in hypersonic flow. Numerical errors introduced by the irregularities of the grid near the shock create non-physical variations that convect downstream to the boundary layer and corrupt surface heat transfer predictions [98].

As will be demonstrated in this thesis, discretizations using higher-order elements can effectively eliminate the errors introduced by an unstructured mesh that is poorly aligned with a shock. Furthermore, when combined with output-based adaptation, automated and accurate aerothermal simulations of hypersonic flows can be realized.

1.2 Thesis Objective

The objective of this work is to develop a robust shock capturing scheme for an adaptive, higher-order discontinuous Galerkin finite element method and apply it to model problems in supersonic and hypersonic flows.

1.3 Background

1.3.1 Higher-Order Methods

The motivation for higher-order discretizations stems from the ability to achieve engineering-required error tolerances with reduced computational load. Finite volume codes are the industry standard approach to CFD for compressible, shock-dominated flows. Higher-order methods are not commonplace in the finite volume community, despite the significant growth in computational resources. Instead, second-order accurate finite volume codes are the most prevalent. Higher-order spatial accuracy for finite volumes typically requires polynomial reconstruction of cell or nodal averages. This creates an expanded numerical stencil, which in turn complicates boundary condition discretizations and adversely impacts iterative algorithms.

Recently, the discontinuous Galerkin (DG) finite element method (FEM) has become a viable alternative to finite volume schemes on unstructured meshes. In the DG context, higher-order approximations are realized by increasing the order of the approximating polynomial, p , within each element. This serves to maintain a nearest neighbor numerical stencil for all solution orders at the cost of additional unknowns to be solved on a given mesh. Thus, in a DG formulation, element-wise coupling only arises via the flux at the

discontinuous element boundaries. The compactness of the DG FEM scheme makes it well suited for parallelization, unstructured grids, and adaptation. Higher-order DG solutions ($p > 1$), for subsonic flows, have shown that strict error tolerances can be achieved for outputs of engineering interest with many fewer degrees of freedom than standard, second-order accurate methods [42, 101].

DG methods were first introduced by Reed and Hill [118] for the neutron transport equation. Much later, the ground work for DG methods applied to non-linear hyperbolic problems was laid down by Karniadakis, Cockburn, and Shu [28–31, 33–35, 124, 141]. Independently, Allmaras and Giles [3, 4] developed a second-order DG scheme for the Euler equations, building off of the work of van Leer [132–135]. Bassi and Rebay and Bey and Oden notably demonstrated the capabilities of DG for both the Euler and Navier-Stokes equations (including Reynolds Averaged Navier-Stokes) [14–16, 18, 19]. Recent work has also focused on improving DG solution methods [44, 88, 95, 107, 140].

1.3.2 Shock Capturing

Discontinuities exist in the solution of many hyperbolic conservation laws. Shocks and contact discontinuities can manifest themselves in the solution to scalar equations, such as Burgers’ equation, or a system of conservation equations, such as the Euler equations which govern inviscid fluid flow. Numerical schemes designed to solve these partial differential equations (PDEs) must be able to capture any discontinuity that might arise in the solution.

The key ingredient for shock capturing in numerical schemes is dissipation. One can think of the numerical solution as an inexact solution to the exact governing equation or, alternatively, as an exact solution to an inexact governing equation [78]. Meaning, the discrete approximation generated by the numerical scheme is an exact solution to a slightly perturbed partial differential equation, called the *modified equation*. The modified equation contains second, third, or higher-order derivatives of the state variable(s). For first-order solutions, or monotone schemes, (where the errors in the solution decrease by $\mathcal{O}(h)$, h being a measure of grid size), the truncation error contains second-order derivatives in the state variables. In the modified equation, these second-order terms have dissipative effects on the numerical solution, leading to smooth numerical solutions. Unfortunately, this creates too much dissipation and smears out discontinuities. In contrast, higher-order accurate schemes have too little numerical dissipation. In fact, many higher-order discretizations have third derivatives in the modified equations, which causes *dispersion*, a phase error for higher frequency modes. Meaning, the speed of wave propagation depends on the wave number itself. Therefore, since a discontinuity contains energy at all wave numbers, the dispersive properties of the numerical scheme will cause oscillations focused at the discontinuity. The errors can also spread to smooth flow regions and corrupt solution accuracy on a global level. Hence, a tradeoff exists between arriving at a physically plausible solution with poor accuracy, and arriving at a higher-order solution in smooth flow regimes with non-physical fluctuations caused by discontinuities.

The resolution of the dichotomy between poor accuracy and non-physical solutions is achieved through the non-linear addition of dissipation via *shock capturing*. Shock capturing involves the use of numerical damping on the higher-order solution to remove the oscillations

near discontinuities. A wide spectrum of approaches to effect this damping exist, some of which predate the advent of the modern computer. The prevalent shock capturing methods in DG are based upon approaches that have been previously used in the continuous finite element, finite volume, and finite difference communities with good success.

Shock Fitting

One alternative to shock capturing is *shock fitting*. Shock fitting involves determination of the shock location within the computational domain through analytical, experimental, or numerical means. The shock location is treated as a boundary condition of sorts within the computational analysis and higher-order accuracy can be attained away from the shock. While this might ostensibly appear as an attractive alternative to crafting intricate shock capturing capabilities for a numerical scheme, it is often not a pragmatic approach. Computational analysis is almost always performed to simulate an unknown flow field, so locating shocks *a priori* can be difficult. Additionally, 3D shock topology can be quite complex, where shocks can bifurcate or end inside of a cell.

Limiters

One of the older and more successful classes of shock capturing methods is the Total Variation Diminishing (TVD) approach. Bounded total variation in a scheme implies that no new local extrema are created, the values of local minima do not decrease, and the values of local maxima do not increase [83]. TVD schemes are generally classified as either flux limiters or slope limiters, the latter of which has become one of the more popular approaches to shock capturing in DG. Slope limiting originated in a series of papers by van Leer [132–135] and focuses on reducing the gradients in a cell based on the values of the neighboring cells so that the solution becomes TVD. For DG, Cockburn and Shu developed a scheme with Runge-Kutta time stepping and a slope limiter based on the minmod operator [28, 30–32, 34]. This method, commonly referred to as RKDG, has a simple implementation, making it both popular [119, 122] and amenable to customization [22]. In flow regions where the limiter is active, the approximating polynomial is reduced to a piecewise-constant representation, leading to a solution that is total variation diminishing in the mean values of each element (TVDM). Unfortunately, the RKDG implementation has some inherent disadvantages, such as the difficulty in marching the residual to a steady-state solution. Specifically, since limiting is applied outside of the residual calculation, the solution that satisfies a zero steady-state residual has spurious oscillations in it.

Reconstruction Methods

Instead of reducing the polynomial order near discontinuities, an alternative approach is to retain the higher-order modes and utilize the additional degrees of freedom to yield sharper shock transitions. The post-TVD generation of shock-capturing schemes produced methods known as *essentially non-oscillatory* (ENO), and were first proposed by Harten et al. [60, 61] and later refined by Shu and Osher [125]. The ENO method chooses a stencil to reconstruct a higher-order polynomial representation from a set of local cell average

values while eliminating spurious oscillations. The ENO schemes, based on their simplicity, sharp shock transitions, and arbitrary orders of accuracy, became quite popular [41]. An improvement over the traditional ENO scheme is the weighted essentially non-oscillatory limiter (WENO) [86]. WENO uses multiple candidate stencils, non-linearly weighted by the smoothness of the solution, whereas ENO adaptively chooses only a single stencil.

A few researchers have also applied the ENO class of shock capturing schemes to DG formulations. Since standard WENO reconstructions require large candidate stencils, Qiu and Shu [113, 114] developed a WENO scheme using Hermite polynomials (HWENO) to maintain the compact DG stencil and demonstrated results on structured 1D and 2D meshes. Compactness is achieved by using the derivatives of the solution, which are readily available in DG FEM, in the reconstruction. The size of the stencil required to achieve a given level of accuracy is therefore smaller than standard WENO methods where the derivatives are not used. Luo et al. [89] advanced this work further and implemented the HWENO scheme on unstructured meshes in 2D and 3D. Unfortunately, the polynomial reconstruction methods of both Qiu and Shu and Luo et al. also occur outside of the residual evaluation, and, similar to the RKDG scheme, obstruct the use of implicit time stepping techniques. However, implicit WENO schemes have been developed in the finite volume community [143, 144], and it is possible that compact, implicit HWENO methods might soon appear in DG as well.

Artificial Viscosity

As mentioned above, some amount of additional dissipation must be added into higher-order numerical schemes to avoid spurious oscillations. One approach is to explicitly add in this additional dissipation in the region of discontinuities by introducing viscous terms to the governing partial differential equation. Viscosity that is on the order of the resolution length scale of the discretization smears out discontinuities until they can be well represented. To ensure consistency of the numerical approximation, this *artificial viscosity* must disappear as $h \rightarrow 0$ and not impact the solution in smooth flow regimes.

The concept of flexible augmentation of artificial viscosity based upon the nature of the solution originated in 1950 by von Neumann and Richtmyer [139]. It was also notably adopted by Baldwin and MacCormack [10] and Jameson et al. [74]. This approach has long been the preferred method of shock capturing in the context of streamwise upwind Petrov-Galerkin (SUPG) finite element methods, as proposed by Hughes et al. [68–71]. Researchers such as Hartmann and Houston [62, 63] and Aliabadi et al. [2] have adopted this approach for use in DG as well, with good results, albeit only for $p = 1$ polynomial solutions.

Persson and Peraire [108] introduced a p -dependent artificial viscosity and demonstrated that higher-order representations and a piecewise-constant artificial viscosity can be combined to produce sub-cell shock resolution. Specifically, introducing an artificial viscosity that scales with the DG resolution length scale, h/p , makes the shock width also scale in the same manner, $\delta_s = Ch/p$. Thus, for sufficiently high p , as shown in Figure 1-2, the shock can be captured within a single element. To locate the shocks in the flow field, Persson and Peraire developed a sensor based on the magnitude of the highest-order coefficients in an




Figure 1-2: Shock layer resolution in DG FEM for increasing polynomial orders.

orthonormal representation of the solution.

This work builds upon the benefits of artificial viscosity for shock capturing in DG. As will be described in Chapter 3, a non-smooth artificial viscosity, such as that used by Persson and Peraire, has some inherent shortcomings. Specifically, element-to-element variations can lead to oscillations in state gradients and disparate equilibrium shock-jump conditions in neighboring elements. This can potentially corrupt the smoothness and accuracy of the downstream flow field. This thesis develops a smoother representation of artificial viscosity, without sacrificing the compact numerical stencil of DG. This is done by allowing the artificial viscosity to be determined by its own PDE, which is appended to the system of governing equations. Thus, while maintaining compactness, the vector of unknown variables is expanded.

Spectral Viscosity

The p -dependent artificial viscosity for DG described above scales the viscosity by the highest mode number in the discretization. In the vanishing viscosity method for spectral elements, proposed by Tadmor and collaborators [26, 45, 57, 92, 93, 129–131], each mode is affected by a different viscosity coefficient, based upon the wave number. In this approach, artificial viscosity is applied to a selection of the highest modes in the scheme (typically modes greater than the square root of the highest wave number of the discrete solution) with a $1/p$ scaling as well. This ensures that the solution converges and prevents oscillations from corrupting the accuracy in smooth flow regions (although some oscillations might remain near discontinuities). Additional post-processing can recover spectrally accurate solutions [45, 92].

A notable variation of the spectral viscosity method is the multi-scale viscosity approach proposed by Oberai and Wanderer [100]. This technique applies different viscosities to the low and high frequency components, the values of which are determined dynamically by a Germano identity [99]. Brachet [21] succeeded in implementing the multi-scale viscosity

methodology for Burgers' equation in DG as well.

1.3.3 Error Estimation and Grid Adaptation

Mesh Adaptation

Mesh adaptation is a widely used and accepted strategy for improving the accuracy of a computational simulation while limiting the increase in computational cost. There are four general approaches to adjust the degrees of freedom (DOF) of a given mesh. The first is p -adaptation, where the interpolation order is locally modified [128]. While p -adaptation can achieve excellent error convergence for smooth flows, difficulties arise near singularities or discontinuities. This contrasts with the most popular adaptation method, h -adaptation, where the local element size is modified. When combined with unstructured and anisotropic mesh generation capabilities, h -adaptation can improve mesh efficiency in boundary layers, wakes, shocks and near complex geometries. A related method, r -adaptation, is a simpler variation of h -adaptation. Instead of regenerating a new mesh at every adaptation iteration, r -adaptation moves node locations without changing the mesh topology to improve the solution accuracy. The final approach is hp -adaptation, where adjustments in h and p are used in conjunction with one another. In this setting h -adaptation is employed for non-smooth flow regions in the vicinity of singularities, and p -adaptation is used in smooth flow regions. Sometimes the choice of adaptation strategy in a particular element, h and/or p , is unclear and criteria must be developed to aid that decision [67].

Error Estimation

The utility of adaptation can be greatly improved if the process is automated. Mesh adaptation based upon user input can be time intensive and requires previous CFD experience. Extricating the user from the adaptation loop is possible if a local estimation of the error can be automatically generated for any given flow solution. Rigorous error estimation can also convey to the user the fidelity level of a computational simulation and allow for informed management of risk and uncertainty in engineering analysis or design.

Purely local measures of error can lead to false confidence in engineering outputs for convection-dominated flows. For instance, local error estimates tied to grid adaptation might lead to considerable grid refinement near shocks or separated boundary layers in transonic flow. However, perturbations in the upstream flow field can have significant impact on the actual shock or separation location, which can, in turn, dramatically change lift and drag predictions.

Estimating the error in an output functional captures the propagation effects inherent in convection-dominated flows by incorporating the adjoint, the solution to the dual problem. The adjoint is a powerful tool that relates local errors to the output and is commonly used in the error analysis for functional outputs and in the calculation of variable sensitivities for gradient-based design optimization. The dual problem, however, requires linearization of the governing PDE and a functional output. For flows with shocks and other singularities, this linearization might not be accurate. The insufficient regularity of the solution might therefore interdict the use of adjoint analysis.

Lindquist and Giles [85] investigated linearized perturbations of unsteady pressure forces in a quasi-1D setting with shocks. They determined that if the shock was smeared enough by artificial viscosity to be resolved by the discretization, then the solution was sufficiently regular to give accurate lift perturbation estimates. Giles [47] later studied the adjoint solution of the 1D, unsteady Burgers equation and found that if there was insufficient numerical dissipation in the scheme, then the adjoint approximation would not converge to the analytic solution. However, if the shocks were regularized with numerical dissipation, then the adjoint solution converged to the analytic adjoint distribution. Pierce and Giles [109] continued this investigation further and determined that accurate error estimates and corrections of output functionals for shocked flows could be obtained if the shocks were smeared by artificial viscosity.

Babuska and Miller [7, 8, 9] were perhaps the first to rigorously frame the error estimate of a numerical solution in terms of a functional quantity of interest. Using FEM solutions of structural analysis problems, they recast outputs such as point stresses and displacements, as integral quantities. The error in the numerical approximation of the outputs could then be expressed in terms of the finite element solution and the adjoint solution. Later, Machiels et al. [91] computed upper and lower bounds for functional outputs of an FEM simulation. The method required primal and dual solutions on a coarse or working mesh as well as a fine mesh solution where discretization errors were negligible. The functional bounds could also be divided into elemental contributions and serve as a guide for automated adaptation.

Becker and Rannacher [17] are responsible for the development of the *dual-weighted residual* method, the approach adopted in this work. They borrowed from the duality techniques in optimal control and exploited the inherent orthogonality of Galerkin FEM to estimate errors for functional outputs. By multiplying local residuals with sensitivity weights, the adjoint solution, they were able to obtain asymptotically correct error estimates. Becker and Rannacher applied the dual-weighted residual method to both linear and non-linear problems and also used it as feedback in an adaptation loop. Much work has been done by others to apply the dual-weighted residual technique to the DG variant of FEM with minor implementation differences [43, 63, 66, 87, 127]. This extension to DG includes demonstrations on non-linear systems of conservation laws, such as the Navier-Stokes equations, and non-linear output functionals as well.

A number of researchers have extended the dual-weighted residual method to other discretizations, such as finite difference and finite volume schemes where no Galerkin orthogonality exists [12, 46, 137]. Barth and Larson [12] estimated the output error for finite volume methods by performing a higher-order reconstruction of the piecewise-constant cell averages. This results in a set of broken polynomials, similar to DG, and facilitates the error estimate. Venditti and Darmofal [136, 137] take a different approach and rely on a fine-mesh solution approximation to anchor the error estimate.

In addition to the ability to estimate the error in a functional output using the adjoint solution, Giles et al. [50, 109] also correct the functional value to achieve greater accuracy. They build on the super-convergent properties of some FEM outputs and apply it to general discretizations, such as finite volume or finite difference. The adjoint-based error correction improves the accuracy of outputs for both linear and non-linear systems, including those

with shocks.

This work is based on the error estimation analysis of Fidkowski [43], who employs a dual-weighted residual error estimate for integral engineering outputs (e.g. drag, lift, far-field pressure etc.). This error estimate is tied to an unstructured, anisotropic, h -adaptation framework at constant p . Where appropriate, minor modifications are made to better support discontinuous flows. Other researchers have successfully demonstrated output-based error estimation in DG via the adjoint for transonic flows using shock capturing with a stabilization method similar to artificial viscosity [62, 63]. This work seeks to advance the capabilities of adaptation for shock flows by applying the framework to higher-order solutions of problems in the supersonic and hypersonic regime.

1.4 Thesis Overview

The primary contributions of this work are the following:

- Motivation for a smooth representation of artificial viscosity for shock capturing in higher-order solutions and a formulation to achieve that representation in the context of the compressible Navier-Stokes equations.
- Modification of dual-weighted residual error estimation and adaptation framework for flows with discontinuities and application to supersonic and hypersonic cases.
- Demonstration of accurate surface heating, shear stress, and pressure prediction for hypersonic problems using unstructured and adapted grids.

Chapter 2 details the DG FEM discretization for convection-diffusion problems and reviews the compressible Navier-Stokes equations. Included in the review is the modification to the governing equations to append an artificial viscosity matrix for shock capturing. Chapter 3 motivates the use of a smooth, higher-order representation of artificial viscosity by highlighting the difficulties of a non-smooth formulation in one and two dimensions. Chapter 4 then presents the chief innovation of this research, a PDE for the control of artificial viscosity, and provides additional comparisons to a non-smooth formulation. Chapter 5 reviews in detail the error estimation and adaptation technique used in this work. In particular, the contribution of artificial viscosity to the error estimate is highlighted. Attention then turns towards the applications mentioned above, the first of which is mesh adaptation for the estimation of drag and pressure signals in a sonic boom problem. Hypersonic applications of the new artificial viscosity model, specifically those focused on the use of unstructured grids and adaptation to predict surface heating, are presented in Chapter 6. Conclusions and areas of future work are summarized in Chapter 7.

Chapter 2

Discontinuous Galerkin Discretization and the Compressible Navier-Stokes Equations

Although the bulk of this work focuses on the compressible Navier-Stokes equations, the shock capturing methodology described in this thesis is applicable to general equation sets in which discontinuities might arise. This chapter presents a general discontinuous Galerkin finite element discretization for nonlinear equations with convective and diffusive terms, including the modifications associated with the addition of artificial viscosity. Additionally, the compressible Navier-Stokes equations and an artificial viscosity matrix for the preservation of total enthalpy are described here.

2.1 Discontinuous Galerkin Finite Elements

Let $\mathbf{u}(\mathbf{x}, t) : \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \mathbb{R}^m$ be the vector of m -state variables in d -dimensions for a general conservation law in the physical domain, $\Omega \subset \mathbb{R}^d \times \mathbb{R}^+$, given in the strong form by,

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathcal{F}(\mathbf{u}) - \nabla \cdot \mathcal{F}^v(\mathbf{u}, \nabla \mathbf{u}) = 0 \quad \text{in } \Omega, \quad (2.1)$$

where $\mathcal{F}(\mathbf{u}) : \mathbb{R}^m \rightarrow \mathbb{R}^{m \times d}$ is the inviscid flux vector and $\mathcal{F}_v(\mathbf{u}, \nabla \mathbf{u}) : \mathbb{R}^m \times \mathbb{R}^{md} \rightarrow \mathbb{R}^{m \times d}$ is the viscous flux.

The discontinuous Galerkin finite element discretization proceeds by deriving a weak form of (2.1). The domain is subdivided by a triangulation, \mathcal{T}_H , into a set of non-overlapping elements, κ , such that $\Omega = \bigcup_{\kappa \in \mathcal{T}_H} \kappa$. Also, define a vector-valued function space of discontinuous, piecewise-polynomials of degree p , \mathcal{V}_H^p , where

$$\mathcal{V}_H^p \equiv \{v \in L^2(\Omega) \mid v|_{\kappa} \in P^p, \forall \kappa \in \mathcal{T}_H\}.$$

The weak form of the governing equations is obtained by multiplying (2.1) by a test function, $\mathbf{v}_H \in (\mathcal{V}_H^p)^m$, and integrating by parts. The solution, $\mathbf{u}_H(\cdot, t) \in (\mathcal{V}_H^p)^m$, satisfies the semi-linear weighted residual (linear in the second argument),

$$\mathcal{R}(\mathbf{u}_H, \mathbf{v}_H) = 0, \quad \forall \mathbf{v}_H \in (\mathcal{V}_H^p)^m,$$

where

$$\mathcal{R}(\mathbf{u}_H, \mathbf{v}_H) \equiv \sum_{\kappa} \left[\int_{\kappa} \mathbf{v}_H^T \frac{\partial \mathbf{u}_H}{\partial t} d\mathbf{x} + \mathbb{E}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H) + \mathbb{V}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H) \right], \quad (2.2)$$

with $\mathbb{E}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H)$ and $\mathbb{V}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H)$ representing the contributions of the inviscid and viscous terms, respectively. Specifically,

$$\mathbb{E}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H) = - \int_{\kappa} \nabla \mathbf{v}_H^T \cdot \mathcal{F}(\mathbf{u}_H) d\mathbf{x} + \int_{\partial\kappa} \mathbf{v}_H^+ \widehat{\mathbf{F}}(\mathbf{u}_H^+, \mathbf{u}_H^-, \hat{\mathbf{n}}) ds,$$

where $\widehat{\mathbf{F}}$ is an approximate flux function, $\hat{\mathbf{n}}$ is the outward pointing normal and the notation $()^+$ and $()^-$ refers to data on the interior and exterior of an element boundary, respectively. Boundary conditions are enforced weakly, by appropriately setting $\widehat{\mathbf{F}}$ when $\partial\kappa$ coincides with $\partial\Omega$. The specific implementation of boundary conditions for the Navier-Stokes equations can be found in Oliver [101] and Fidkowski et al. [44].

The viscous flux contributions are discretized according to the second form of Bassi and Rebay [16] (BR2). In this approach, (2.1) is written as a system of equations,

$$\frac{\partial u}{\partial t} + \nabla \cdot \mathcal{F} - \nabla \cdot \mathcal{Q} = 0 \quad (2.3)$$

$$\mathcal{Q} - \mathbf{A}_v \nabla \mathbf{u} = 0, \quad (2.4)$$

where it is assumed that \mathcal{F}_v has a linear dependence on the state gradients, $\mathcal{F}_v(\mathbf{u}, \nabla \mathbf{u}) = \mathbf{A}_v(\mathbf{u}) \nabla \mathbf{u}$ and $\mathbf{A}_v \in \mathbb{R}^{md \times md}$ is the viscosity matrix.¹ The first equation is multiplied by a test function, $\mathbf{v}_H \in (\mathcal{V}_H^p)^m$, and second equation is multiplied by a test function, $\boldsymbol{\tau}_H \in (\mathcal{V}_H^p)^{md}$. After an integration by parts, one obtains,

$$\sum_{\kappa} \left[\int_{\kappa} \mathbf{v}_H^T \frac{\partial \mathbf{u}_H}{\partial t} d\mathbf{x} + \mathbb{E}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H) + \int_{\kappa} \nabla \mathbf{v}_H^T \cdot \mathcal{Q}_H d\mathbf{x} - \int_{\partial\kappa} \mathbf{v}_H^+ \widehat{\mathcal{Q}} \cdot \hat{\mathbf{n}} ds \right] = 0 \quad (2.5)$$

$$\sum_{\kappa} \left[\int_{\kappa} \boldsymbol{\tau}_H^T \cdot \mathcal{Q}_H d\mathbf{x} + \int_{\kappa} \mathbf{u}_H^T \nabla \cdot (\mathbf{A}_v^T \boldsymbol{\tau}_H) d\mathbf{x} - \int_{\partial\kappa} (\widehat{\mathbf{A}_v \mathbf{u}})^T \boldsymbol{\tau}_H^+ \cdot \hat{\mathbf{n}} ds \right] = 0 \quad (2.6)$$

where $\widehat{(\cdot)}$ denotes a numerical flux approximation for discontinuous data.

The last two terms in (2.5) are the viscous contributions, $\mathbb{V}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H)$ in (2.2). They can be further manipulated by letting $\boldsymbol{\tau}_H \equiv \nabla \mathbf{v}_H$ in (2.6) and substituting the first term in (2.6) into the third term in (2.5). One more integration by parts yields the general

¹Some implications of non-linear dependence upon the state gradients are addressed in Appendix A

Table 2.1: Viscous fluxes

| | $\widehat{\mathcal{Q}}$ | $\widehat{\mathbf{A}_v \mathbf{u}}$ |
|---------------------|---|-------------------------------------|
| Interior | $\{\mathbf{A}_v \nabla \mathbf{u}_H\} - \eta_f \{\boldsymbol{\delta}_f\}$ | $\mathbf{A}_v^+ \{\mathbf{u}_H\}$ |
| Boundary, Dirichlet | $\mathbf{A}_v^b \nabla \mathbf{u}_H^+ - \eta_f^b \boldsymbol{\delta}_f^b$ | $\mathbf{A}_v^b \mathbf{u}^b$ |
| Boundary, Neumann | $\mathbf{A}_v^b \nabla \mathbf{u}^b$ | $\mathbf{A}_v^+ \mathbf{u}_H^+$ |

discretization of diffusion terms,

$$\mathbb{V}_\kappa(\mathbf{u}_H, \mathbf{v}_H) = \int_\kappa \nabla \mathbf{v}_H^T \cdot \mathbf{A}_v \nabla \mathbf{u}_H d\mathbf{x} + \int_{\partial\kappa} \nabla(\mathbf{v}_H^T)^+ \left(\widehat{\mathbf{A}_v \mathbf{u}} - \mathbf{A}_v^+ \mathbf{u}_H^+ \right) \cdot \hat{\mathbf{n}} ds - \int_{\partial\kappa} \mathbf{v}_H^+ \widehat{\mathcal{Q}} \cdot \hat{\mathbf{n}} ds. \quad (2.7)$$

The next steps involve choosing numerical approximations for the fluxes, $\widehat{\mathbf{A}_v \mathbf{u}}$ and $\widehat{\mathcal{Q}}$. While there are a number of choices that lead to consistent discretizations, not all of these options lead to stable, compact, and dual consistent schemes [5]. As mentioned above, the results in this work were obtained with the BR2 discretization. In this context, the choices for $\widehat{\mathbf{A}_v \mathbf{u}}$ and $\widehat{\mathcal{Q}}$ are described in Table 2.1, where $\{\cdot\}$ denotes the average operator on an element boundary, $\{w\} = 0.5[w^+ + w^-]$, the superscript, b , refers to data from an appropriately constructed boundary state, the subscript, f , refers to a given face, η is a stabilization parameter, and $\boldsymbol{\delta}_f, \boldsymbol{\delta}_f^b \in (\mathcal{V}_H^p)^{md}$ are auxiliary variable components for interior and boundary faces. These are defined such that, $\forall \boldsymbol{\tau}_H \in (\mathcal{V}_H^p)^{md}$.

$$\begin{aligned} \int_{\kappa^+} \boldsymbol{\tau}_H^T \cdot \boldsymbol{\delta}_f^+ d\mathbf{x} + \int_{\kappa^-} \boldsymbol{\tau}_H^T \cdot \boldsymbol{\delta}_f^- d\mathbf{x} &= \int_{\sigma_f} [\![\mathbf{u}_H]\!]^T \cdot \{\mathbf{A}_v^T \boldsymbol{\tau}_H\} ds \\ \int_\kappa \boldsymbol{\tau}_H^T \cdot \boldsymbol{\delta}_f^b d\mathbf{x} &= \int_{\sigma_f^b} (\mathbf{u}_H^+ - \mathbf{u}^b)^T [\mathbf{A}_v^T \boldsymbol{\tau}_H \cdot \hat{\mathbf{n}}]^b ds \end{aligned}$$

where $[\![\cdot]\!]$ is the jump-operator on an element boundary, $[w] = w^+ \hat{\mathbf{n}}^+ + w^- \hat{\mathbf{n}}^-$, σ_f and σ_f^b are interior and boundary faces, respectively, with κ^+ and κ^- denoting elements on either side of σ_f .

2.1.1 Solution and Geometry Interpolation

The function space, \mathcal{V}_H^p , consists of discontinuous, piecewise-polynomials. This work used a polynomial basis for \mathcal{V}_H^p such that the discrete solution could be written as a linear combination of basis functions,

$$\mathbf{u}_H(\mathbf{x}) = \sum_k \mathbf{U}_{H_k} \phi_k(\boldsymbol{\xi}(\mathbf{x})),$$

where $\{\phi\}$ is the set of basis functions locally supported on a single element and \mathbf{U}_H is the discrete solution vector. Here the basis functions are defined on a canonical reference element in reference space, $\boldsymbol{\xi} \in \mathbb{R}^d$. Employing a reference element, a triangle in two dimensions and a tetrahedron in three dimensions, allowed for simple use of existing quadrature rules for integral evaluation and also facilitated the use of high-order, curved elements. For

high-order elements, the additional nodes were equally spaced within the reference element and corresponded to given positions in \mathbf{x} . The reference-to-global mapping was,

$$\mathbf{x} = \sum_k \mathbf{x}_k \phi_k(\boldsymbol{\xi}),$$

where \mathbf{x} is the global coordinate, $\boldsymbol{\xi}$ is the reference coordinate, ϕ_k is the Lagrange basis function associated with node n_k , and \mathbf{x}_k is the global coordinate of that same node. In the Lagrange basis, the value of ϕ_i on node n_j is given by the Kronecker delta, δ_{ij} , where the nodes are evenly spaced within the reference element.

2.1.2 Solution Method

Although the focus of this work was on steady-state solutions, the unsteady term was retained to improve the initial transient behavior of the solver. Specifically, backward Euler time stepping was used such that the discrete solution vector, \mathbf{U}_H , at time interval, $n + 1$, is given by,

$$\mathbf{U}_H^{n+1} = \mathbf{U}_H^n - \left[\frac{1}{\Delta t} \mathcal{M}_H + \frac{\partial \mathbf{R}_H}{\partial \mathbf{U}_H} \right]^{-1} \mathbf{R}_H(\mathbf{U}_H^n) \quad (2.8)$$

where \mathcal{M}_H is the mass matrix and $\mathbf{R}_H(\mathbf{U}_H)$ is the discrete spatial residual vector. To accelerate convergence, especially when the initial condition was far from the steady-state solution, the time step, Δt , was incrementally increased.

The solution of (2.8) requires the inversion of the Jacobian matrix. Given the large size of the Jacobian in DG, iterative methods were used to solve the linear system. The results presented here were obtained with the restarted GMRES algorithm. To further aid iterative convergence of the linear system, an ILU factorization is used as a preconditioner where the factorization is performed using a reordering of elements into lines [39]. The lines are a unique set of elements created by the coupling between elements in a $p = 0$ discretization of a scalar, linear convection-diffusion equation [44].

It should also be noted that all of the higher-order solutions presented in this thesis were arrived at via p -sequencing. Meaning, lower-order solutions served as the initial condition for higher-order solutions. This was found to be a robust path towards higher-order solutions, especially for large Mach and Reynolds numbers. For grid adaptation though, once an order p solution was obtained on an initial mesh using p -sequencing, low order solutions were no longer necessary for flow initialization on the later adapted meshes. Instead, the flow was initialized by transferring the order p solution to the next adapted mesh. Details on the solution transfer for adaptation are found in Section 5.2.

2.2 Compressible Navier-Stokes Equations

The compressible Navier-Stokes equations are a non-linear system that can be written in the form of (2.1). In this context, the conservative state vector is, $\mathbf{u} = [\rho, \rho v_i, \rho E]^T$, where ρ is the density, v_i is the velocity in the i -th coordinate direction and E is the total internal

energy. The inviscid flux vector is, using index notation,

$$\mathcal{F}_i(\mathbf{u}) = \begin{bmatrix} \rho v_i \\ \rho v_i v_j + \delta_{ij} p \\ \rho v_i H \end{bmatrix},$$

where p is the static pressure, $H = E + p/\rho$ is the total enthalpy, and δ_{ij} is the Kronecker delta. The pressure is related to the state vector by the equation of state,

$$p = (\gamma - 1)\rho \left(E - \frac{1}{2} v_i v_i \right),$$

where γ is the ratio of specific heats ($\gamma = 1.4$ in air).

The viscous flux vector, $\mathcal{F}^v = \mathbf{A}_v \nabla \mathbf{u}$, can similarly be written using index notation as,

$$\mathcal{F}_i^v(\mathbf{u}, \nabla \mathbf{u}) = \begin{bmatrix} 0 \\ \tau_{ij} \\ v_j \tau_{ij} + \kappa_T \frac{\partial T}{\partial x_i} \end{bmatrix},$$

where τ is the shear stress defined below, κ_T is the thermal conductivity, $T = p/\rho R$ is the temperature and R is the gas constant. The shear stress is,

$$\tau_{ij} = \mu \left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} \right) - \delta_{ij} \lambda \frac{\partial v_k}{\partial x_k},$$

where μ is the dynamic viscosity and $\lambda = -\frac{2}{3}\mu$ is the bulk viscosity coefficient. Here the dynamic viscosity is assumed to adhere to Sutherland's Law,

$$\mu = \mu_{ref} \left(\frac{T}{T_{ref}} \right)^{1.5} \frac{T_{ref} + T_s}{T + T_s},$$

and the thermal conductivity is related to the viscosity by the Prandtl number, Pr ,

$$\kappa_T = \frac{\gamma \mu R}{(\gamma - 1) Pr}.$$

In air, $T_{ref} = 288^\circ \text{K}$ (unless the freestream value is noted), $T_s = 110.4^\circ \text{K}$, and $Pr = 0.71$.

2.3 Artificial Viscosity Matrix

When artificial viscosity is added to the system for the purposes of resolving discontinuities, the viscous flux is augmented such that $\mathcal{F}^v = (\mathbf{A}_v + \mathbf{A}_\epsilon) \nabla \mathbf{u}$, where $\mathbf{A}_\epsilon(\mathbf{u}) : \mathbb{R}^m \rightarrow \mathbb{R}^{md \times md}$

is an anisotropic, diagonal viscosity matrix defined as,

$$\mathbf{A}_\epsilon \nabla \mathbf{u} = \hat{\epsilon}(\mathbf{u}, \mathbf{h}) \text{diag} \left[\mathcal{I} \frac{h_i}{\bar{h}} \right] \nabla \mathbf{u}, \quad (2.9)$$

$$\hat{\epsilon} \equiv \begin{cases} 0, & \epsilon \leq \theta_L \\ \frac{1}{2} \theta_H \left(\sin \left[\pi \left(\frac{\epsilon - \theta_L}{\theta_H - \theta_L} - \frac{1}{2} \right) \right] + 1 \right), & \theta_L < \epsilon < \theta_H \\ \theta_H, & \epsilon \geq \theta_H \end{cases}$$

where $\hat{\epsilon} : \mathbb{R}^m \times \mathbb{R}^d \rightarrow \mathbb{R}$ is the artificial viscosity applied, \bar{h} is the arithmetic mean of, $\mathbf{h}(\mathbf{x}) \in \mathbb{R}^d$, a local vector-measure of the element size described below, and $\mathcal{I} \in \mathbb{R}^m$ is a vector of ones. $\hat{\epsilon}$ is scaled to smoothly vary between zero and a maximum value, θ_H , as ϵ , the artificial viscosity produced by the shock capturing method, varies between θ_L , a minimum value, and θ_H . The determination of ϵ , based on a non-linear shock switch, will be described in greater detail in Chapter 4. For consistency, θ_L and θ_H scale with $\lambda \bar{h}/p$ ($\theta_L = 0.01 \lambda_{\max} \bar{h}/p$ and $\theta_H = \lambda_{\max} \bar{h}/p$), and λ_{\max} is the maximum wave speed of the system.

The local measure of element size is a linear variation throughout the computational mesh. Using continuous, linear, nodal basis functions, $\mathbf{h}(\mathbf{x})$ can be written as,

$$\mathbf{h}(\mathbf{x}) = \sum_{k=1}^{d+1} \mathbf{H}_k \varphi_k(\boldsymbol{\xi}(\mathbf{x})),$$

where $\mathbf{H}_k \in \mathbb{R}^d$ is the average value of the bounding box vectors of all elements bordering the k -th principal node of an element and φ_k is the nodal, linear basis function associated with the node. The arithmetic mean, $\bar{h}(\mathbf{x})$, is therefore a continuously varying scalar function throughout the domain,

$$\bar{h}(\mathbf{x}) = \frac{1}{d} \sum_{i=1}^d \sum_{k=1}^{d+1} H_{ki} \varphi_k(\boldsymbol{\xi}(\mathbf{x})).$$

2.3.1 Numerical Diffusion for Constant Total Enthalpy

The addition of artificial viscosity via (2.9) is valid for all systems of equations in the form of (2.1). For compressible flow, the artificial viscosity matrix can be modified to better preserve the behavior of the shock transition given by the Euler equations. The Rankine-Hugoniot shock jump relations state that, for the steady Euler equations, total enthalpy is conserved across the shock [77]. For unsteady flow cases, there is no such guarantee of constant total enthalpy in the flowfield. Since the focus of this thesis is on steady-state solutions, when dealing with the compressible Navier-Stokes equations, $\mathbf{A}_\epsilon(\mathbf{u})$ in (2.9) is substituted with $\tilde{\mathbf{A}}_\epsilon(\mathbf{u}) : \mathbb{R}^m \rightarrow \mathbb{R}^{md \times md}$, an artificial viscosity matrix designed to preserve total enthalpy.

Isenthalpic formulations of the Euler equations have long been considered in the computational community. Lytton [90] and Jameson [72] are two examples of numerical discretizations designed to preserve total enthalpy throughout the flow field. In the steady Euler equations, the energy and mass equations are identical if the constant factor, H , is

removed from the energy equation. However, discrete schemes do not necessarily satisfy this property. A discrete solution with constant H is admissible if the numerical dissipation applied to the energy equation reduces to the numerical dissipation applied to the continuity equation when ρH is replaced by ρ [75]. Consequently, the application of the artificial viscosity matrix to the conservative state vector according to (2.9) would violate this criteria because the dissipation in the energy equation would act on ρE . Thus, $\tilde{\mathbf{A}}_\epsilon$ is defined by,

$$\tilde{\mathbf{A}}_\epsilon \nabla \mathbf{u} = \mathbf{A}_\epsilon \nabla \tilde{\mathbf{u}}, \quad (2.10)$$

where $\tilde{\mathbf{u}} = [\rho, \rho v_i, \rho H]^T$.

Another formulation of an artificial viscosity matrix for the preservation of total enthalpy uses the Navier-Stokes viscosity matrix for an ideal gas. If the Prandtl number is set to, $Pr = 0.75$, one can show that this choice gives shock transitions with constant total enthalpy. This approach is used by Persson and Peraire [108], but is not applied to the results presented in this work.

The Numerical Flux Function

In addition to the artificial viscosity matrix, numerical diffusion is added to the DG FEM scheme through the approximate flux function as well. For a flux function to ensure that total enthalpy is constant throughout the domain, it must also apply the same dissipation to the energy equation as the continuity equation multiplied by the constant factor, H . Jameson [75] describes and presents a few approximate flux functions that satisfy this criteria. In this work, two different approximate flux functions were employed. The first is the Roe flux function [120], which does not ensure constant total enthalpy, and the other is the modification of the van Leer flux difference splitting by Hänel [59], which is designed to preserve total enthalpy.

Chapter 3

Motivation for Smooth Artificial Viscosity

The contributions of this thesis work revolve around a new model for artificial viscosity that produces a smooth variation of viscosity from one element to the next. To motivate this new model, this chapter highlights some of the shortcomings of a non-smooth representation of viscosity, as applied to higher-order DG interpolations. The comparison is made between a piecewise-constant and smooth variation of viscosity in both a one-dimensional and multi-dimensional setting. Before continuing with the comparison, though, a discussion on the use of artificial viscosity for shock capturing is presented.

3.1 Vanishing Viscosity and Conservation Laws

Consider a scalar, non-linear, hyperbolic conservation law in one spatial dimension,

$$\begin{aligned} u_t + \mathcal{F}(u)_x &= 0, & \text{in } \Omega, \\ u &= 0, & \text{on } \partial\Omega, \\ u(x, 0) &= u_0(x) \end{aligned} \tag{3.1}$$

where $u(x, t) : \mathbb{R} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ is the state variable of interest, $\mathcal{F}(u) : \mathbb{R} \rightarrow \mathbb{R}$ is the flux function and the domain is $\Omega \subset \mathbb{R} \times \mathbb{R}^+$. This conservation law relates derivatives of the solution in time to derivatives of the flux function in space. Independent of the smoothness of the initial condition, the solution may develop discontinuities, in which case these derivatives become undefined. It is therefore more convenient to seek weak solutions, where (3.1) is multiplied by a test function, $\phi \in C_0^1(\mathbb{R} \times \mathbb{R}^+)$, and integrated by parts to shift derivatives from $u(x, t)$ to $\phi(x, t)$,

$$\iint_{\Omega} [\phi_t u + \phi_x \mathcal{F}(u)] dx dt = 0 \quad . \tag{3.2}$$

Unfortunately, solutions to (3.2) are not necessarily unique. Additional constraints are required to single out the physically relevant solution. This additional constraint is called the *entropy condition*, and imposes a companion conservation law on the weak solution. The additional conservation law involves a convex function, $\eta(u)$, which is called *entropy*.

When combined with an *entropy flux*, $\psi(u)$, the entropy should be conserved in smooth flows,

$$\eta(u)_t + \psi(u)_x = 0 \quad . \quad (3.3)$$

Comparing the linearized form of (3.3),

$$\eta'(u)u_t + \psi'(u)u_x = 0 \quad ,$$

with the linearized form of (3.1) multiplied by $\eta'(u)$,

$$\eta'(u)u_t + \eta'(u)\mathcal{F}'(u)u_x = 0 \quad ,$$

it follows that $\psi'(u) = \eta'(u)\mathcal{F}'(u)$, or

$$\psi(u) = \int_0^u \eta'(\xi)\mathcal{F}'(\xi)d\xi \quad .$$

The linearization, $\eta'(u)$, yields the *entropy variables*, which have some appealing mathematical properties for hyperbolic systems [11].

For smooth flows, the physically admissible weak solution is the one that satisfies the entropy condition. This solution is also the *vanishing viscosity* limit of the original conservation law [38],

$$\lim_{\nu \rightarrow 0} [u_t + \mathcal{F}(u)_x = \nu u_{xx}] \quad . \quad (3.4)$$

LeVeque [83] demonstrates the link between the vanishing viscosity solution and the entropy condition by multiplying (3.4) by $\eta'(u)$,

$$\begin{aligned} \lim_{\nu \rightarrow 0} \eta'(u)u_t + \eta'(u)\mathcal{F}'(u)u_x &= \eta'(u)\nu u_{xx} \\ \eta(u)_t + \psi(u)_x &= \nu \left[(\eta'(u)u_x)_x - \eta''(u)u_x^2 \right] \quad , \end{aligned}$$

and integrating over a space-time slab, $[x_1, x_2] \times [t_1, t_2]$, to yield,

$$\begin{aligned} \lim_{\nu \rightarrow 0} \int_{t_1}^{t_2} \int_{x_1}^{x_2} \eta(u)_t + \psi(u)_x \, dx \, dt &= \\ \nu \int_{t_1}^{t_2} \left[\eta'(u(x_2, t))u_x(x_2, t) - \eta'(u(x_1, t))u_x(x_1, t) \right] dt &- \nu \int_{t_1}^{t_2} \int_{x_1}^{x_2} \eta''(u)u_x^2 \, dx \, dt. \end{aligned}$$

For smooth flow, as $\nu \rightarrow 0$, the right-hand-side of the equation tends towards zero, verifying the link between the vanishing viscosity solution and the entropy condition. However, for discontinuous flows (the motivating factor for the use of weak solutions), the first term on the right-hand-side is zero, but the second term involves integrations of, u_x^2 , over the space-time slab. If the vanishing viscosity limit involves discontinuities, then u_x is infinite, and this term will not be zero in the limit. However, knowing that $\nu > 0$, $u_x^2 > 0$ and $\eta''(u) > 0$ (by convexity), then

$$\int_{t_1}^{t_2} \int_{x_1}^{x_2} \eta(u)_t + \psi(u)_x \, dx \, dt \leq 0 \quad .$$

This allows for a more general statement of (3.3), called the *entropy inequality*,

$$\eta(u)_t + \psi(u)_x \leq 0 \quad ,$$

which states that the total integral of entropy, $\eta(u)$, can only decrease. This entropy inequality is satisfied in the weak sense, and selects the physically admissible weak solution as the vanishing viscosity limit.

3.1.1 Burgers' Equation Example

To demonstrate that the vanishing viscosity limit of a hyperbolic conservation law for discontinuous flow yields the exact solution, consider the traveling wave problem for Burgers' equation,

$$u_t + \left(\frac{1}{2}u^2\right)_x = 0, \quad u(x, 0) = \begin{cases} u_L, & x < 0 \\ u_R, & x > 0 \end{cases} \quad .$$

The exact solution is given by,

$$u(x, t) = \begin{cases} u_L, & x - st < 0 \\ u_R, & x - st > 0 \end{cases} \quad , \quad (3.5)$$

where s is the shock speed, $s = (u_L + u_R)/2$.

To obtain the vanishing viscosity limit of this problem, begin with the viscous Burgers equation,

$$u_t + \left(\frac{1}{2}u^2\right)_x = \nu u_{xx} \quad ,$$

and a change of variables, $\xi = x - st$, such that the solution is only a function of one variable, $u = u(\xi)$,

$$-su' + uu' = \nu u'' \quad .$$

Integrating once with respect to ξ and applying the boundary conditions for u and u' at $\xi = \pm\infty$ gives,

$$(u - u_L)(u - u_R) = 2\nu u' \quad .$$

The solution can be obtained by solving the integral,

$$\begin{aligned} \int d\xi &= 2\nu \int \frac{du}{(u - u_L)(u - u_R)} \\ \xi &= \frac{2\nu}{u_L - u_R} \ln \left(\frac{u_L - u}{u - u_R} \right) \\ u(\xi) &= \frac{u_R \exp \left[\left(\frac{u_L - u_R}{2\nu} \right) \xi \right] + u_L}{\exp \left[\left(\frac{u_L - u_R}{2\nu} \right) \xi \right] + 1} \quad . \end{aligned}$$

This expression utilizes the fact that $\frac{d}{dx} \ln(-u) = (-\frac{1}{u})(-\frac{du}{dx}) = (\frac{1}{u})(\frac{du}{dx}) = \frac{d}{dx} \ln(u)$ and that $u_L > u > u_R$ to ensure that the logarithm argument is positive. Evaluating the vanishing

viscosity limit of the above expression,

$$\lim_{\substack{\nu \rightarrow 0 \\ \xi \rightarrow 0^-}} u(\xi) = u_L, \quad \lim_{\substack{\nu \rightarrow 0 \\ \xi \rightarrow 0^+}} u(\xi) = u_R.$$

which conforms to the exact solution given by (3.5).

The traveling wave solution, for three different values of viscosity, is plotted in Figure 3-1a. As the viscosity becomes smaller, the traveling wave solution approaches a true discontinuity and the analytic solution of (3.5).

This traveling wave solution also offers some insight as to the connection between the viscosity and the shock width. Define the shock width, δ_s , to be the ξ -distance that the maximum slope of u (which occurs at $\xi = 0$) takes to traverse the difference between u_L and u_R ,

$$\delta_s \equiv \frac{u_R - u_L}{u'(0)} = \frac{8\nu}{u_L - u_R} \quad ,$$

as shown in the diagram of Figure 3-1b. If u_L and u_R are simplified to be $\pm u_0$, respectively, then the shock width is,

$$\delta_s = \frac{4\nu}{u_0} \quad .$$

Thus, the shock width varies directly with the amount of viscosity applied. Additionally, the stronger the shock, the smaller the viscous shock layer. For the purposes of shock capturing, artificial viscosity added to the discretization on the order of the resolution length scale ensures that the shock can be accurately resolved by the numerical scheme.

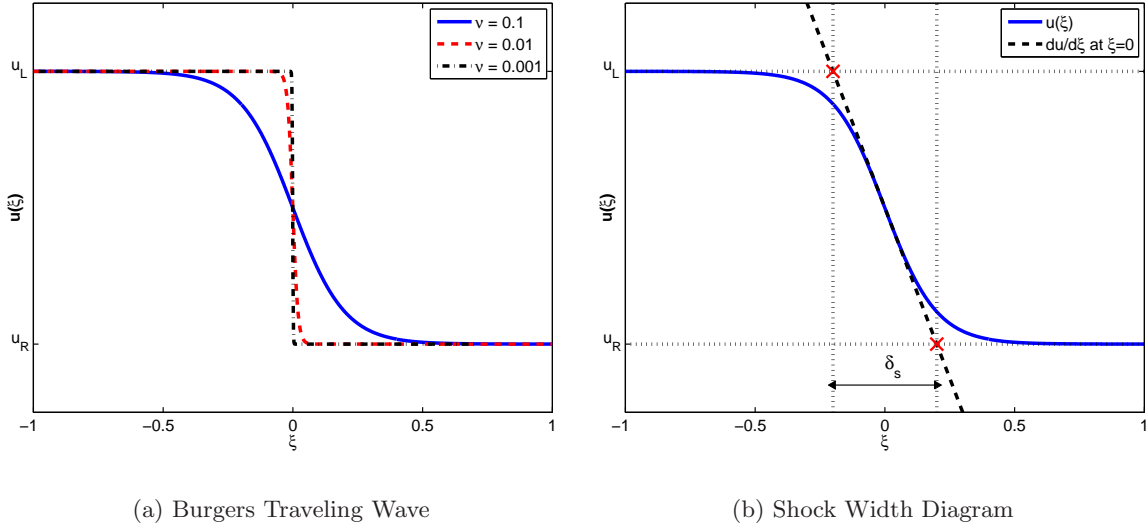


Figure 3-1: Traveling wave solution of Burgers' equation with vanishing viscosity and shock width diagram.




Figure 3-2: Distributions of piecewise-constant and Gaussian artificial viscosity as applied to the 1D modified Burgers equation.

3.2 One-Dimensional Results

The one-dimensional viscous Burgers equation is employed to demonstrate the benefits of a smooth variation in artificial viscosity, compared to a non-smooth representation. The governing equation is modified to support a steady-state shock solution,

$$\frac{\partial}{\partial x} \left(\frac{1}{2} u^2 \right) = \alpha u + \frac{\partial}{\partial x} \left(\nu(x) \frac{\partial u}{\partial x} \right) + f(x), \quad \text{for } x \in \Omega \subset \mathbb{R} \quad (3.6)$$

where $u(x, t)$ is the state variable, $\nu(x)$ is the viscosity, α is a constant ($\alpha = -0.1$) and the forcing term, $f(x)$, is set such that the exact, steady-state solution with $\nu = 0$ has a shock at $x = 0$,

$$u(x) = \begin{cases} 2 + \sin\left(\frac{\pi x}{2}\right), & x < 0 \\ -2 - \sin\left(\frac{\pi x}{2}\right), & x > 0 \end{cases} \quad (3.7)$$

The viscosity, $\nu(x)$, is prescribed to be either a piecewise-constant or smooth Gaussian function, as depicted in Figure 3-2. The piecewise-constant viscosity is applied to the cells immediately adjacent to the shock location with adjustable amplitude. The Gaussian distribution of viscosity is specified to have a standard deviation equal to the cell size and the same total area as the piecewise-constant rectangle between $x \in [-h, h]$.

To perform the comparison, (3.6) is discretized using sixth order Legendre polynomials as the basis functions in DG FEM. The 1D, scalar Lax-Friedrich's flux function,

$$\widehat{\mathbf{F}}(u^+, u^-) = \frac{1}{2} \left([\mathcal{F}(u^+) + \mathcal{F}(u^-)] - C[u^- - u^+] \right); \quad C = \max_{s \in [u^-, u^+]} |\mathcal{F}'(s)|,$$

is used for the jumps in the inviscid fluxes and the BR2 discretization described in the previous Chapter, with the numerical fluxes listed in Table 2.1, is used to discretize the viscous fluxes. A global L^2 -norm of the solution error and an H^1 -norm of the error outside of the shock layer is measured between the discrete and exact solution for the two viscosity formulations. The error norms are defined as,

$$\begin{aligned}\|u - u_H\|_{L^2} &= \left[\int_{\Omega} (u - u_H)^2 dx \right]^{1/2} \\ \|u - u_H\|_{H^1} &= \left[\int_{\Omega/\delta_s} (u - u_H)^2 + \left(\frac{du}{dx} - \frac{du_H}{dx} \right)^2 dx \right]^{1/2},\end{aligned}$$

where u_H is the discrete solution and δ_s is the shock layer, defined to be the distance extending from $x = 0$ to where the discrete solution is first within 0.5% of the exact solution. While the norms can take on different values depending on the mesh length scale, the domain size and triangulation were held fixed for all cases.

The comparison between the two viscosity representations is made at three different viscosity amplitudes. The results for low, moderate, and high values of viscosity are shown in Figure 3-3 with the error norm comparison in Table 3.1. At a low viscosity amplitude, the numerical oscillations in $u(x)$ are damped, but oscillations still remain in the derivative, $u_x(x)$, for both solutions, suggesting that the shock is not entirely resolved. At a higher viscosity amplitude, the Gaussian viscosity solution is smooth for both $u(x)$ and $u_x(x)$, but the piecewise-constant viscosity solution still has significant oscillations in $u_x(x)$. These oscillations are due to the conservation of the flux, $(u^2/2 + \nu u_x)$, across element boundaries. A jump in ν , requires a similar jump in u_x , but there is no change in u_x that can compensate for a jump to $\nu = 0$. Thus, for higher-order solutions, this jump in the viscosity induces derivative fluctuations throughout the element. Finally, for much higher viscosity amplitudes, the Gaussian viscosity solution remains well-behaved, but the piecewise-constant solution suffers from oscillations in both $u(x)$ and $u_x(x)$.

Since the solutions for $u(x)$ are quite similar, except at the highest viscosity amplitude, the L^2 -norm values of the solution error are also quite similar. However, the greater accuracy of the Gaussian viscosity solution is reflected in the H^1 -norm values. At the high viscosity amplitude, the H^1 -norm of the error for the Gaussian viscosity solution is smaller than the piecewise-constant solution by two orders of magnitude.

To achieve a smooth variation of viscosity, the Gaussian distribution in Figure 3-2 has a larger viscosity footprint than the piecewise-constant representation. To emphasize that the driver of the oscillations in u_x observed in Figure 3-3 are due to a jump in ν to zero, and not the larger viscosity footprint, the same study is performed with an expanded distribution of piecewise-constant viscosity, as shown in Figure 3-4. In this application, the piecewise-constant representation of viscosity is decreased in a staircase manner in the four cells immediately adjacent to the shock, $x \in [-2h, 2h]$.

The results of the expanded piecewise-constant representation of viscosity, shown in Figure 3-5, are similar to those in Figure 3-3. Despite the larger viscosity footprint, the oscillations in u_x for the piecewise-constant viscosity solution do not dissipate at any vis-

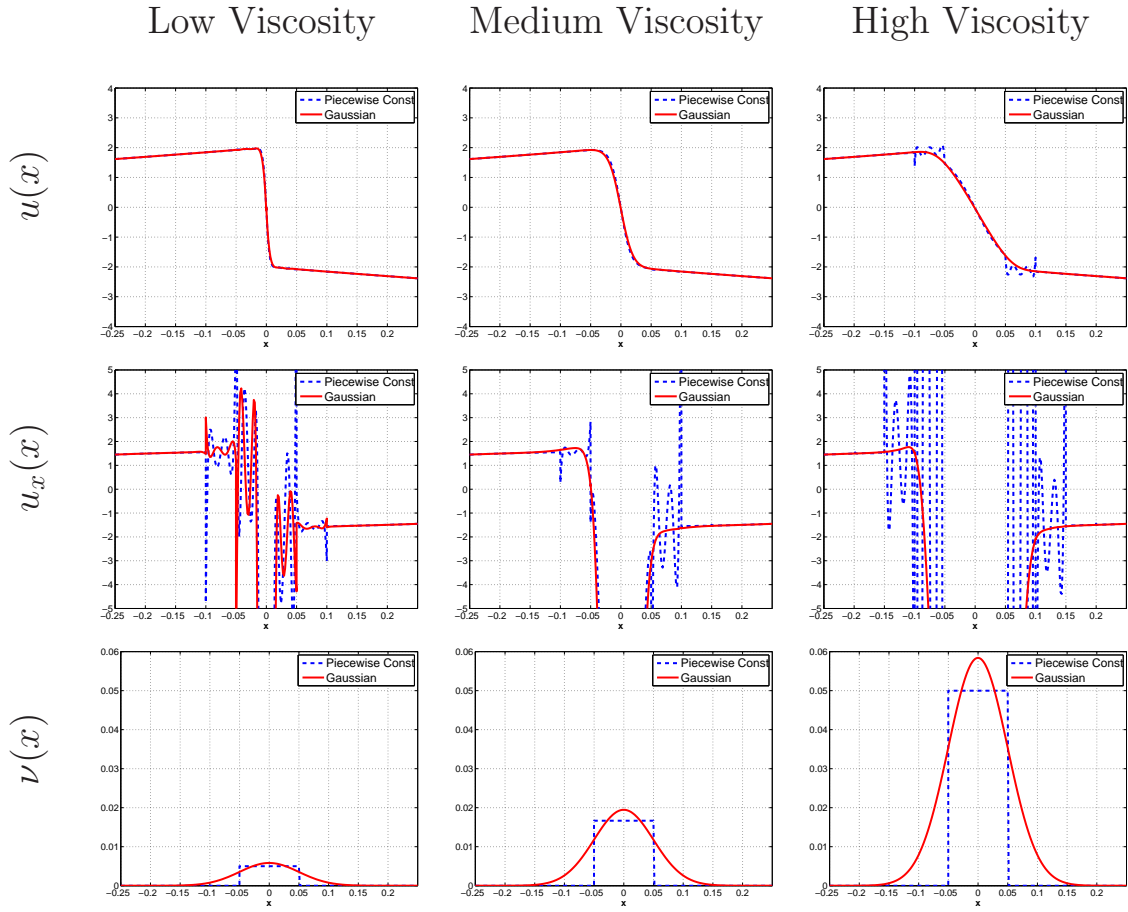


Figure 3-3: Comparison of piecewise-constant and Gaussian viscosity solutions for modified Burgers equation across three different viscosity amplitudes (40 elements, $p = 6$).

Table 3.1: Global L^2 error norm and H^1 error norm outside of shock layer comparisons of piecewise-constant, expanded piecewise-constant, and Gaussian viscosity solutions for modified Burgers equation across three different viscosity amplitudes (40 elements, $p = 6$)

| | Low Viscosity | | Medium Viscosity | | High Viscosity | |
|----------------------------|---------------|-------------|------------------|-------------|----------------|-------------|
| | H^1 -norm | L^2 -norm | H^1 -norm | L^2 -norm | H^1 -norm | L^2 -norm |
| Piecewise-const | 0.873 | 0.124 | 0.680 | 0.227 | 19.196 | 0.387 |
| Exp Piecewise-const | 0.852 | 0.124 | 0.364 | 0.227 | 0.875 | 0.389 |
| Gaussian | 0.548 | 0.134 | 0.180 | 0.244 | 0.167 | 0.405 |




Figure 3-4: Distributions of expanded piecewise-constant and Gaussian artificial viscosity as applied to the 1D modified Burgers equation.

cosity amplitude. Furthermore, the oscillations in u_x occur when $\nu = 0$ at $x = \pm 2h$, while the viscosity jump at $x = \pm h$ introduces a slope discontinuity but not oscillations throughout the element. At the highest viscosity amplitude, however, the expanded application of the piecewise-constant viscosity has removed the oscillations in u observed in Figure 3-3. The H^1 -norms of the error for the expanded piecewise-constant solution in Table 3.1 are also much smaller than the original values, albeit still significantly larger than those of the Gaussian solution.

3.3 Adjoint Analyses

The shock capturing methodology in this work is intended to be well-suited to adjoint-based analyses. Specifically, the adjoint-based analyses of interest are design variable sensitivity in the context of gradient-based optimization and error estimates of output functionals.

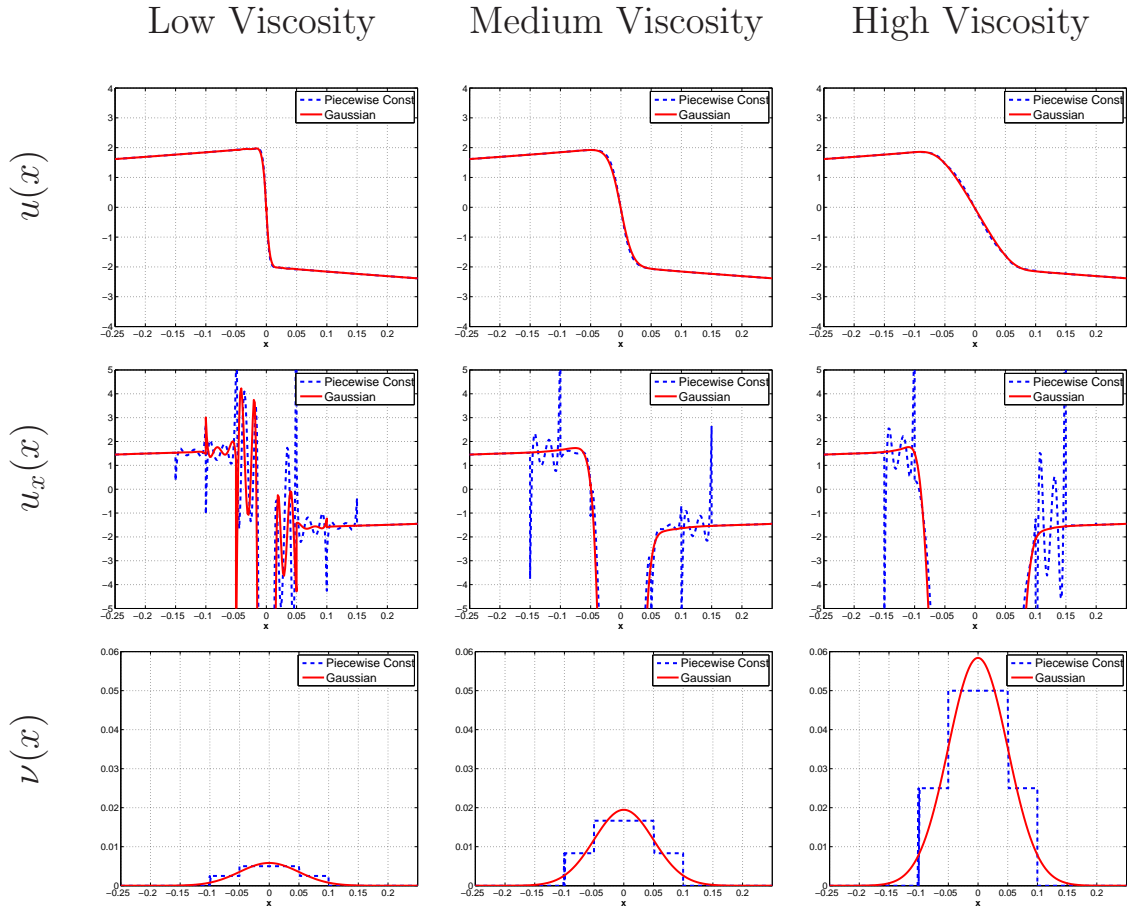


Figure 3-5: Comparison of expanded piecewise-constant and Gaussian viscosity solutions for modified Burgers equation across three different viscosity amplitudes (40 elements, $p = 6$).

Error estimation is discussed at length in Chapter 5. The use of adjoint-based approach to obtain design variable sensitivities, specifically for cases that involve discontinuities, is discussed here.

The adjoint can be understood as a Green's function that relates a source of a PDE to a functional output computed from the PDE solution. For the purposes of computational error estimation, the source is the discretization error of the numerical scheme. In gradient-based design optimization, the source is the perturbation due to a small change in the design variables.

To observe the connection between the adjoint and design variable sensitivities, let $\beta \in \mathbb{R}^k$ be a vector of design variables and $\mathbf{u} \in \mathbb{R}^n$ be the discrete solution to a governing PDE,

$$\mathcal{R}(\beta; \mathbf{u}) = 0.$$

The perturbations in the residual due to perturbations in the design vector lead to,

$$\frac{\partial \mathcal{R}}{\partial \mathbf{u}} \frac{d\mathbf{u}}{d\beta} = -\frac{\partial \mathcal{R}}{\partial \beta},$$

where $\partial \mathcal{R} / \partial \mathbf{u} \in \mathbb{R}^{n \times n}$ is the Jacobian matrix. The sensitivity of a functional output, $\mathcal{J}(\beta; \mathbf{u})$, to the design vector can be written as,

$$\frac{d\mathcal{J}}{d\beta} = \frac{\partial \mathcal{J}}{\partial \beta} + \frac{\partial \mathcal{J}}{\partial \mathbf{u}} \frac{d\mathbf{u}}{d\beta} \quad .$$

The adjoint, $\psi \in \mathbb{R}^n$, is defined as,

$$\frac{d\mathcal{J}}{d\beta} = \frac{\partial \mathcal{J}}{\partial \beta} - \psi^T \frac{\partial \mathcal{R}}{\partial \beta} \quad , \tag{3.8}$$

which implies that ψ is the solution to the following $n \times n$ linear system,

$$\frac{\partial \mathcal{R}}{\partial \mathbf{u}}^T \psi = \frac{\partial \mathcal{J}}{\partial \mathbf{u}}^T \quad . \tag{3.9}$$

If the interest is in error estimation, then the design vector may be ignored and instead consider perturbations to the residual due to perturbations in the solution, $\delta \mathbf{u}$,

$$\begin{aligned} \delta \mathcal{R} &= \mathcal{R}(\mathbf{u} + \delta \mathbf{u}) \approx \frac{\partial \mathcal{R}}{\partial \mathbf{u}} \delta \mathbf{u} \\ \delta \mathbf{u} &= \left(\frac{\partial \mathcal{R}}{\partial \mathbf{u}} \right)^{-1} \delta \mathcal{R} \quad . \end{aligned}$$

Variations to the functional are,

$$\delta \mathcal{J} = \mathcal{J}(\mathbf{u} + \delta \mathbf{u}) - \mathcal{J}(\mathbf{u}) \approx \frac{\partial \mathcal{J}}{\partial \mathbf{u}} \delta \mathbf{u}.$$

Applying (3.9), the change in \mathcal{J} can be expressed in terms of the adjoint,

$$\delta\mathcal{J} = \frac{\partial\mathcal{J}}{\partial\mathbf{u}}\delta\mathbf{u} = \boldsymbol{\psi}^T \frac{\partial\mathcal{R}}{\partial\mathbf{u}}\delta\mathbf{u} = \boldsymbol{\psi}^T \delta\mathcal{R}.$$

3.3.1 Design Variable Sensitivity Error

As described in Section 1.3.3, the investigations of Lindquist and Giles [85] showed that variable sensitivities could be obtained in shock-flow cases with regularized, viscous shocks. Giles [48] later demonstrated that for the quasi-1D Euler equations, second-order accurate estimates for the integrated pressure force could be obtained even for flows with shocks. In a design setting, if the number of constraints is small compared to the number of design variables, the computational cost of calculating design variable sensitivities for gradient-based optimization can be reduced with the use of the adjoint. Giles [47] demonstrated that the correct adjoint solution could be obtained with shocks regularized with artificial viscosity. For extensive details regarding the adjoint-approach to design variable sensitivities see the work of Giles and Pierce [49] and Jameson [73].

For the problem described by (3.6), the forcing term was modified to support an exact solution with a design variable, β , while maintaining a shock at $x = 0$,

$$u(x) = \begin{cases} \beta[2 + \sin(\frac{\pi x}{2})], & x < 0 \\ -\beta[2 + \sin(\frac{\pi x}{2})], & x > 0 \end{cases} \quad (3.10)$$

Additionally, an output functional was prescribed to be,

$$\mathcal{J}(u) = \left[\int_{\Omega} (u - 1)^2 dx \right]^{1/2}. \quad (3.11)$$

With these definitions, the error between the computed variable sensitivity via (3.8) and the analytic sensitivity (obtained by differentiating (3.11) directly using the exact solution given in (3.10)) was evaluated for both of the viscosity representations of Figure 3-2. The accuracy study was performed at multiple viscosity amplitudes, over a series of grid refinements and polynomial orders ($p = 1-4$). For the given output, the error in the sensitivity estimates is expected to be $\mathcal{O}(h/p)$.

The sensitivity errors for both viscosity representations, shown in Figure 3-6, agree reasonably well with one another and the assumed rate of $\mathcal{O}(h/p)$. The errors shown have been scaled by their respective values of p , such that all of the lines should ideally lie on top of one another. At the low and high values of viscosity, the sensitivity errors of the piecewise-constant viscosity solution lie noticeably further from the assumed $\mathcal{O}(h/p)$ rate than those of the Gaussian viscosity solution. Nevertheless, the differences in sensitivity error between the two viscosity representations are not as dramatic as those in Figure 3-3. This is not surprising as the output of interest in (3.11) is quite similar to an L^2 -norm, which is well behaved for both viscosity representations.

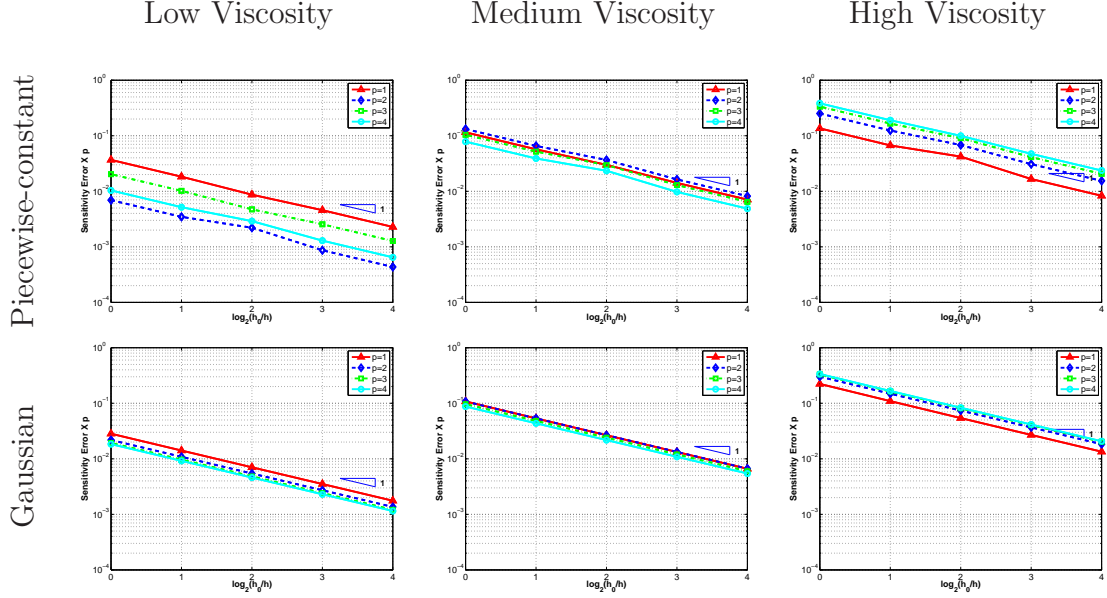




Figure 3-6: Grid convergence study of variable sensitivity errors computed via the adjoint for piecewise-constant and Gaussian distributions of viscosity.

3.4 Multiple Dimension Issues

In one dimension, the errors induced by the shock for the non-smooth viscosity solution are generally confined to a region near the shock. In multiple dimensions, however, this is no longer the case. The jumps in viscosity from one element to the next along a discontinuity, due to changes in the shock strength, cell size and orientation, introduce errors both normal and tangential to the shock front. These errors create entropy gradients in the flow field which can convect downstream and pollute the solution accuracy. This might lead to erroneous surface pressure or heat transfer on a body behind a shock. For instance, consider the inviscid, supersonic flow over a cylinder depicted in Figure 3-7. If a non-smooth representation of artificial viscosity is used to capture the shock, then significant oscillations in total pressure arise behind the shock front. While total pressure behind the bow shock is not constant, it should vary smoothly as the strength of the shock changes due to its curvature. The three lines shown in Figure 3-7b are total pressure measurements taken from three grids across two uniform refinements along the solid black line in Figure 3-7a. As the grid becomes finer, the total pressure oscillations persist and the wavelength decreases. Although the smaller wavelength makes the solution appear to deteriorate in accuracy with refinement, results in Section 4.4.1 show that the global error decreases at the expected convergence rate. Nevertheless, the considerable noise in the total pressure reflects a great deal of uncertainty associated with engineering outputs using non-smooth artificial viscosity. These variations in the solution downstream of a shock were previously observed by Quattrocchi [115].



(a) Mach contours



(b) Total pressure

Figure 3-7: Mach number contours and total pressure measurements along a line behind the bow shock across two grid refinements for a $p = 3$ solution of a 2D flow around a cylinder, $M_\infty = 4$.

Chapter 4

A PDE-Based Artificial Viscosity Model

Using a one-dimensional problem with a known analytic solution, the previous chapter demonstrated the benefits of a smooth variation of artificial viscosity, compared to a non-smooth representation. However, the shock location for a given flow field is rarely known *a priori* and the artificial viscosity cannot be a pre-determined function in space. Similar to the example problem, two artificial viscosity models were explored in this research: a non-smooth and a higher-order polynomial representation. This chapter describes these artificial viscosity models, and, in particular, a PDE-based viscosity model is presented as a means to achieve a smooth variation in artificial viscosity.

4.1 Non-Smooth Artificial Viscosity

A non-smooth formulation of artificial viscosity can be obtained by allowing the viscosity to be controlled by a non-linear shock switch,

$$\epsilon|_{\kappa} = \frac{\bar{h}(\mathbf{x})}{p} \lambda_{\max}(\mathbf{u}) S_{\kappa}(\mathbf{u}), \quad (4.1)$$

where $S_{\kappa}(\mathbf{u}) : \mathbb{R}^m \rightarrow \mathbb{R}$ is the non-linear switch or indicator function that detects the spurious numerical oscillations in element, κ , and determines the amount of artificial viscosity to add. The specifics of the shock switch are discussed in greater length in Section 4.3. It is important to note here though that the shock switch formulations used for this work are element-based integrals and ϵ therefore is a non-smooth function in the domain with jumps at element boundaries. As proposed by Persson and Peraire [108], the shock switch is multiplied by an h/p -scaling to allow for sub-cell shock resolution.

4.2 PDE-based Artificial Viscosity

The results presented in the previous chapter suggest that a smooth representation of artificial viscosity, without large jumps at element boundaries, offers benefits compared to

a non-smooth viscosity representation. If (4.1) is considered a baseline artificial viscosity model, then the simplest approaches to achieving a smooth variation of $\epsilon(\mathbf{x})$ might be to either use a switch based on purely local quantities or to construct a higher-order patch over non-smooth values. Unfortunately, a pointwise switch, based on purely local quantities, is not yet a dependable option at higher-order interpolations due to the severe numerical noise in the shock layer. Also, a higher-order patch over the shock switch values would extend the numerical stencil in DG because the residual evaluation on an element edge would depend on the elements sharing that edge as well as second degree neighbors. To maintain a compact stencil, a PDE-based model of artificial viscosity is proposed. The drawback to this approach is that additional degrees of freedom are introduced for the artificial viscosity.

The PDE model for artificial viscosity satisfies a diffusion equation of the following form,

$$\frac{\partial \epsilon}{\partial t} = \nabla \cdot \left(\frac{\boldsymbol{\eta}}{\tau} \nabla \epsilon \right) + \frac{1}{\tau} \left[\frac{\bar{h}(\mathbf{x})}{p} \lambda_{\max}(\mathbf{u}) S_{\kappa}(\mathbf{u}) - \epsilon \right] \quad \text{in } \Omega \subset \mathbb{R}^d \times \mathbb{R}^+ \quad (4.2)$$

$$\frac{\partial \epsilon}{\partial \hat{\mathbf{n}}} = \frac{\epsilon_{\infty} - \epsilon}{L} \quad \text{on } \partial\Omega \quad (4.3)$$

where λ_{\max} is the maximum wave speed of the system, τ is an appropriate time constant and $\boldsymbol{\eta} \in \mathbb{R}^{d \times d}$ is the conductivity. The working variable of the PDE is $\epsilon(\mathbf{x}, t) : \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \mathbb{R}$, which has units of kinematic viscosity ($Length^2/Time$), and is also an additional state variable that is appended to the state vector. For the sake of brevity, this PDE will be referred to as the *artificial viscosity equation* and its formulation of ϵ will be referred to as *PDE-based artificial viscosity*.

No physical basis exists that prescribes the boundary conditions for ϵ . Since a shock could conceivably intersect any boundary at any angle, a boundary condition was sought that did not impose any restrictions on the distribution of artificial viscosity. In this light, neither a Dirichlet or Neumann boundary condition for ϵ are appropriate. A homogeneous Dirichlet condition ($\epsilon = 0$) implicitly assumes that the shock terminates at the boundary and a homogeneous Neumann condition assumes that the shock is normal to the boundary. Thus, a Robin boundary condition is most appropriate, where the gradient of ϵ is proportional to the difference between the boundary value and an ambient state ($\epsilon_{\infty} = 0$) over a local length scale, L ($L = 10\mathbf{h} \cdot \hat{\mathbf{n}}$). Additional investigation into the boundary behavior of the artificial viscosity equation is described in the next subsection.

The PDE model for artificial viscosity is designed to address the shortcomings of the non-smooth approach. The shock indicator acts as a forcing term that drives ϵ to be non-zero in the vicinity of discontinuities. As in (4.1), the shock switch is multiplied by the h/p -scaling to allow for sub-cell shock resolution. Also, even though the shock indicator, $S_{\kappa}(\mathbf{u})$, might be an element-integral quantity, higher-order representations of ϵ are still possible. The diffusion term, governed by the parameter, $\boldsymbol{\eta}$, ensures that the viscosity is smooth (no large jumps at element edges) and that artificial viscosity produced in one element diffuses to its neighbors.

The artificial viscosity equation is cast with a time derivative and time constant, τ , defined such that ϵ evolves at least as fast as the primary system of equations. This time scale is chosen to approximate the time it takes the fastest wave speed to traverse the

resolution scale of the solution. In this way, there is no lag between the need for stabilization and the build-up of artificial viscosity. To this end, the time constant, τ , is defined by

$$\tau = \frac{\min_i h_i}{C_1 p \lambda_{\max}(\mathbf{u})}; \quad C_1 = 3.$$

The second parameter of the artificial viscosity equation is the conduction coefficient, $\boldsymbol{\eta} \in \mathbb{R}^{d \times d}$. For dimensional consistency it must have units of $Length^2$, and, so that the viscosity only spreads to neighboring elements, $\boldsymbol{\eta}$ should be made an explicit function of $\mathbf{h}(\mathbf{x})$. Thus, an appropriate setting of $\boldsymbol{\eta}$ is simply,

$$\boldsymbol{\eta} = C_2 \text{diag}([h_x^2, h_y^2, h_z^2]^T),$$

The quantity, $\boldsymbol{\eta}/\tau$, is therefore,

$$\frac{\boldsymbol{\eta}}{\tau} = C_1 C_2 \frac{p \lambda_{\max}(\mathbf{u})}{\min_i h_i} \text{diag}([h_x^2, h_y^2, h_z^2]^T); \quad C_1 C_2 = 15. \quad (4.4)$$

4.2.1 Modified System of Equations

The artificial viscosity equation is an additional PDE that must be solved with the original governing equations. The working variable, ϵ , is appended to the state vector and a source term vector, \mathcal{G} , is included in the system. Thus, (2.1), becomes

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathcal{F}(\mathbf{u}) - \nabla \cdot \mathcal{F}^v(\mathbf{u}, \nabla \mathbf{u}) = \mathcal{G} \quad \text{in } \Omega.$$

The semi-linear weighted residual also includes the source term vector,

$$\mathcal{R}(\mathbf{u}_H, \mathbf{v}_H) \equiv \sum_{\kappa} \left[\int_{\kappa} \mathbf{v}_H^T \frac{\partial \mathbf{u}_H}{\partial t} d\mathbf{x} + \mathbb{E}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H) + \mathbb{V}_{\kappa}(\mathbf{u}_H, \mathbf{v}_H) - \int_{\kappa} \mathbf{v}_H^T \mathcal{G} d\mathbf{x} \right],$$

For the compressible Navier-Stokes equations, the state vector becomes $\mathbf{u} = [\rho, \rho v_i, \rho E, \epsilon]^T$ and the flux and source term vectors are modified to be,

$$\mathcal{F}_i(\mathbf{u}) = \begin{bmatrix} \rho v_i \\ \rho v_i v_j + \delta_{ij} p \\ \rho v_i H \\ 0 \end{bmatrix}, \quad \mathcal{G} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \frac{1}{\tau} \left(\frac{\bar{h}}{p} \lambda S_{\kappa} - \epsilon \right) \end{bmatrix}.$$

The viscous flux includes the terms from the Navier-Stokes equations and the artificial viscosity,

$$\mathcal{F}_i^v(\mathbf{u}, \nabla \mathbf{u}) = \begin{bmatrix} \hat{\epsilon} \frac{h_i}{h} \frac{\partial \rho}{\partial x_i} \\ \tau_{ij} + \hat{\epsilon} \frac{h_i}{h} \frac{\partial \rho v_j}{\partial x_i} \\ v_j \tau_{ij} + \kappa T \frac{\partial T}{\partial x_i} + \hat{\epsilon} \frac{h_i}{h} \frac{\partial \rho H}{\partial x_i} \\ C_1 C_2 \frac{p \lambda h_i^2}{\min_i h_i} \frac{\partial \epsilon}{\partial x_i} \end{bmatrix},$$

4.2.2 Green's Function Behavior

A Green's function analysis of the artificial viscosity PDE offers insight into its behavior for a source impulse and various boundary conditions.

Let $u(x) \in \mathcal{V}$ solve a linear differential equation of the form,

$$\begin{aligned}\mathcal{L}u &= f & \text{in } \Omega \subset \mathbb{R}^d \\ \mathcal{D}u &= 0 & \text{on } \partial\Omega,\end{aligned}$$

where $\mathcal{L} : \mathcal{V} \rightarrow \mathbb{R}$ and $\mathcal{D} : \mathcal{V} \rightarrow \mathbb{R}$ are differential operators. A Green's function, $G(x, s) \in \mathcal{V}$, for the PDE is defined by

$$\begin{aligned}\mathcal{L}G &= \delta(x - s) & \text{in } \Omega \\ \mathcal{D}G &= 0 & \text{on } \partial\Omega,\end{aligned}\tag{4.5}$$

where $\delta(x)$ is the Dirac delta function. A Green's function is therefore constructed as the impulse response for the differential operator. In this way the solution can be expressed as,

$$u(x) = \int_{\Omega} G(x, s) f(s) ds.$$

To investigate the behavior of the artificial viscosity equation, consider a suitable model PDE,

$$\eta^2 \nabla^2 u - u = -f \quad \text{in } \Omega\tag{4.6}$$

The analysis will be performed in 1D. The first case considered is an infinite domain, $\Omega = [-\infty, \infty]$, with $u \rightarrow 0$ as $|x| \rightarrow \infty$. The Green's function solution is,

$$G(x, s) = \begin{cases} \frac{\eta}{2} \exp\left(\frac{x-s}{\eta}\right), & x < s \\ \frac{\eta}{2} \exp\left(\frac{s-x}{\eta}\right), & x > s \end{cases}.\tag{4.7}$$

Figure 4-1a plots the Green's function of the 1D model problem for an impulse at $s(x) = 0$ across various values of η . There are a few features of this plot to highlight. First, the solution is smooth everywhere, except for $x = s$. Second, the solution is positive for positive impulse inputs and the solution decays to zero away from the source. Finally, for increasing values of η , the spreading of $G(x, s)$ also increases due to the increasing quantities of diffusion. Since the value of η determines the peak solution value of the impulse, Figure 4-1b plots the same curves, but scaled so that they all have a peak value of unity.

The results in Figure 4-1 can also be used to select the appropriate level of diffusion for the artificial viscosity equation. A shock in a given element might induce fluctuations in neighboring elements as well. Thus, artificial viscosity engendered in one element should spread such that it has a non-zero value approximately 2-3 elements away. If one cell is interpreted as $h = \Delta x = 1$ (and no change of variables is necessary), then an appropriate value of η is such that $G(x, s)$ has a sizable magnitude at $x = \pm 3$. For $\eta = 1$, the magnitude of $G(\pm 3, 0)$ is quite small while if $\eta = 3$ the value is quite large. Thus, while the criteria used to select η are somewhat subjective, from Figure 4-1b, an appropriate value is somewhere

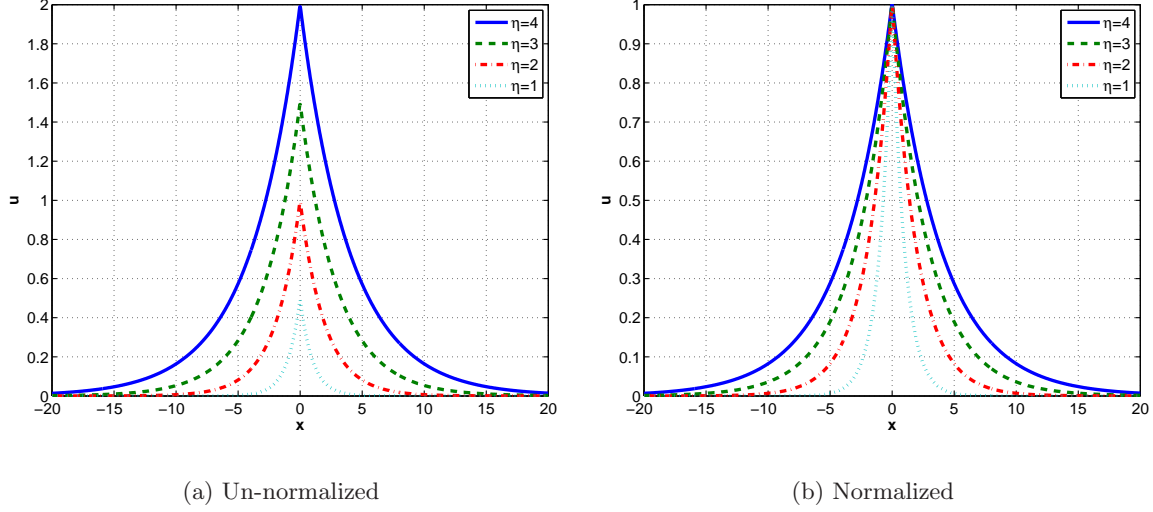


Figure 4-1: Green's function for the 1D model of the artificial viscosity equation.

around $\eta = 2$. This is consistent with the selected value of $C_2 = \eta^2$ in (4.4).

The Green's function expressed in (4.7) is derived from (4.6) with Dirichlet boundary conditions and the results in Figure 4-1 reflect the behavior of the artificial viscosity equation far away from the boundary. To investigate the impact of boundary conditions, additional Green's functions were obtained for a boundary imposed at $x = -a$. For three different boundary conditions (Dirichlet, Neumann and Robin), the Green's function expressions are,

$$u = 0 \text{ at } x = -a : \quad G(x, s) = \begin{cases} \frac{\eta}{2} \left[\exp\left(\frac{x-s}{\eta}\right) - \exp\left(-\frac{2a+x+s}{\eta}\right) \right], & x < s; \\ \frac{\eta}{2} \left[\exp\left(\frac{s-x}{\eta}\right) - \exp\left(-\frac{2a+x+s}{\eta}\right) \right], & x > s; \end{cases}$$

$$\frac{du}{dx} = 0 \text{ at } x = -a : \quad G(x, s) = \begin{cases} \frac{\eta}{2} \left[\exp\left(\frac{x-s}{\eta}\right) + \exp\left(-\frac{2a+x+s}{\eta}\right) \right], & x < s; \\ \frac{\eta}{2} \left[\exp\left(\frac{s-x}{\eta}\right) + \exp\left(-\frac{2a+x+s}{\eta}\right) \right], & x > s; \end{cases}$$

$$\eta \frac{du}{dx} = (u \overset{0}{\longleftarrow} \infty) \text{ at } x = -a : \quad G(x, s) = \begin{cases} \frac{\eta}{2} \exp\left(\frac{x-s}{\eta}\right), & x < s; \\ \frac{\eta}{2} \exp\left(\frac{s-x}{\eta}\right), & x > s; \end{cases}$$

For a boundary located approximately five cells from an impulse at $s = 0$, $\Omega = [-5, \infty]$, and a fixed value of η ($\eta = 2$), the behavior induced by the different boundary conditions on the solution can be observed in Figure 4-2a. At this proximity, there is some minor variation between the three different solutions since the source dies away in a few cells. However, if the impulse is located approximately one cell away from a domain boundary, $\Omega = [-1, \infty]$, which is depicted in Figure 4-2b, the choice of boundary condition plays a significant role in the solution behavior. Both the Dirichlet and Neumann boundary conditions perturb

the solution behavior significantly from the behavior in an infinite domain. Only the Robin boundary condition allows the solution to behave as if the domain boundary were not present. This is reflected mathematically by the fact that the Green's function for a Robin boundary at $x = -a$ is identical to (4.7).

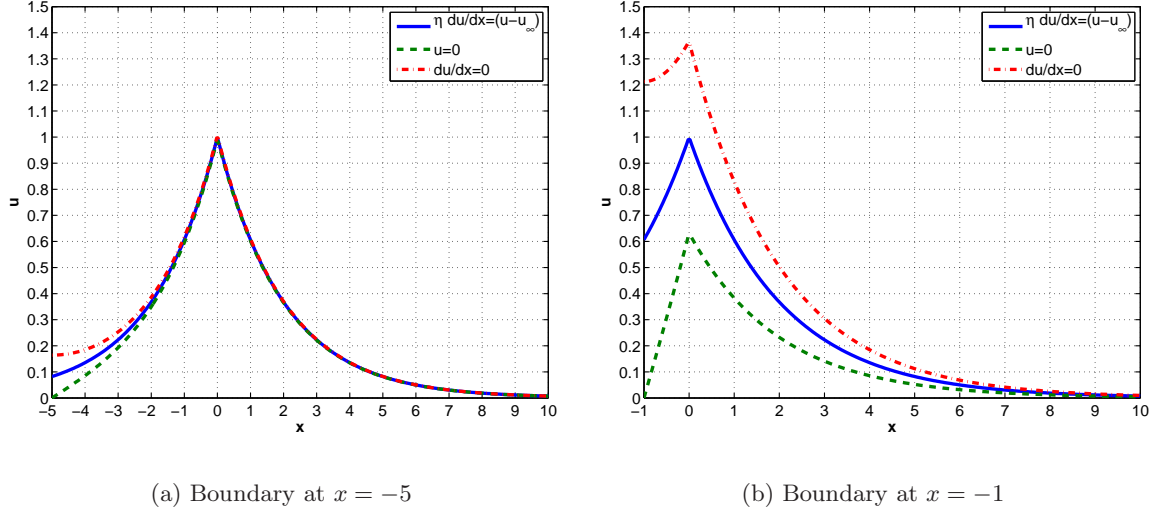


Figure 4-2: Boundary condition impact upon Green's function behavior for a source near the domain boundary.

4.3 Shock Indicators

The shock indicator, $S_\kappa(\mathbf{u})$, can take many forms. This research has employed two different indicators, which are presented in this section. Both of the indicators are element-based integrals leading to a single, scalar measure of the need for dissipation to control the numerical oscillations near a discontinuity.

Many other shock indicators controlling the non-linear addition of artificial viscosity have appeared in the literature. The first to suggest the use of artificial viscosity for shock capturing, von Neumann and Richtmyer [139], used a sensor based on the gradient of the specific volume, which is significantly higher in a shock than in smooth flow. Later, Baldwin and MacCormack [10] and Jameson et al. [74] employed artificial viscosity for their finite volume schemes, added through a sensor based on the second derivative of pressure.

Within the finite element community, the use of artificial viscosity for shock capturing has also been quite popular. In most respects, the variational form of the shock capturing operator takes the form of,

$$\sum_{\kappa} \int_{\kappa} \mathbf{D} \nabla \mathbf{v}_H \cdot \nabla \mathbf{u}_H d\mathbf{x} \quad , \quad (4.8)$$

where \mathbf{D} contains the non-linear switch that controls where dissipation is added and in what quantity. This form originated in a series of papers by, Hughes et al. [68–71], who presented

the streamwise upwind Petrov-Galerkin (SUPG) discretization. For the Euler equations, this included a shock capturing term that Johnson et al. [79] write as,

$$\sum_{\kappa} \bar{\delta} \int_{\kappa} \frac{\left| \frac{\partial \mathbf{u}_H}{\partial t} + \nabla \cdot \mathcal{F}(\mathbf{u}_H) \right|}{\sigma + |\nabla \bar{\mathbf{u}}_H|} \nabla \mathbf{v}_H \cdot \nabla \bar{\mathbf{u}}_H d\mathbf{x} \quad , \quad (4.9)$$

where $\bar{\mathbf{u}}$ denotes entropy variables and $\bar{\delta}$ and σ are $\mathcal{O}(h^\alpha)$ quantities with $\alpha \approx 1$. The shock switch in this case is the strong form of the residual of the Euler equations. For the Navier-Stokes equations, the viscous terms are included in the numerator of (4.9). The residual is large in regions where the flow is not well resolved, such as near discontinuities, and small in smooth flow regions. With this shock sensor, Johnson et al. proved that SUPG FEM solutions to Burgers' equation converge to the entropy solution.

In the discontinuous Galerkin context, the shock capturing term in the form of (4.8), using the residual switch in (4.9), is quite popular with minor implementation differences. Bassi and Rebay [13] implemented a discretization that cast \mathbf{D} as,

$$\mathbf{D} = Ch_\kappa^2 \left[\sum_k \left(\frac{\nabla \cdot \mathcal{F}^k(\mathbf{u}_H)}{|u|_H^k + c} \right)^2 \right]^{-\frac{1}{2}}$$

where C and c are empirical constants and the summation over k is a summation over the rank of the system of equations. Similarly, Hartmann and Houston [63] write \mathbf{D} as,

$$\mathbf{D} = Ch^{2-\beta} |\nabla \cdot \mathcal{F}(\mathbf{u}_H)| \mathbf{I},$$

where C and β are positive constants and $\mathbf{I} \in \mathbb{R}^m$ is the identity matrix. Hartmann [62] later modified the length scale for anisotropic diffusion in anisotropic elements, as is done in this work. The above expressions control the artificial viscosity by some measurement of the strong-form residual of the compressible Euler equations. It should be noted that these expressions use purely local quantities. However, the aforementioned DG researchers used only linear ($p = 1$) basis functions, and, unfortunately, as the order of the polynomial increases, so too does the noise in the shock layer for derivatives of the state variables. Thus, one can no longer use strictly local switches and must instead rely on element-integral quantities.

4.3.1 Resolution Indicator

A resolution-based indicator was introduced by Persson and Peraire [108] as their method of detecting shocks to demonstrate the sub-cell shock capturing capabilities of artificial viscosity with higher-order, DG solutions. This indicator treats the higher-order solution as though it were comprised of a sequence of Fourier modes. For smooth flows, the coefficients of increasing Fourier modes should die away rapidly. In a true discontinuity, however, all frequency modes are present. This idea is similar to error indicators for adaptation in spectral methods [94]. With this concept in mind, the state vector at any point in a

higher-order approximation, can be represented as,

$$u(x) = \sum_k^{N(p)} U_k \phi_k(x),$$

where ϕ_k are the basis functions, U_k are the associated weights and $N(p)$ is the size of the higher-order expansion of degree p . Assuming an orthogonal basis,

$$\hat{u}(x) = \sum_k^{N(p-1)} U_k \phi_k(x),$$

where \hat{u} is the truncated representation of the state vector at order $p - 1$.

With the definitions of u and \hat{u} , the resolution indicator can be defined by,

$$F_\kappa = \log_{10} \left(\frac{\langle f - \hat{f}, f - \hat{f} \rangle}{\langle f, f \rangle} \right) \quad (4.10)$$

where $\langle \cdot, \cdot \rangle$ represents the standard L^2 inner-product, and $f = f(\mathbf{u}) : \mathbb{R}^m \rightarrow \mathbb{R}$ is a component or function of the state vector. As with Persson and Peraire, this work relies on density as a reliable quantity for $f(\mathbf{u})$.

The final scaling of the indicator used by Persson and Peraire is such that it varies smoothly between zero and a maximum value,

$$S_\kappa(F_\kappa; \theta_S, \psi_0, \Delta\psi) = \begin{cases} 0, & F_\kappa \leq \psi_0 - \Delta\psi \\ \theta_S, & F_\kappa \geq \psi_0 + \Delta\psi \\ \frac{\theta_S}{2} \left(1 + \sin \frac{\pi(F_\kappa - \psi_0)}{2\Delta\psi} \right), & |F_\kappa - \psi_0| < \Delta\psi \end{cases}$$

where θ_S is a maximum value ($\theta_S = 1$) and ψ_0 and $\Delta\psi$ are empirically determined constants. In a Fourier expansion, coefficients decay at the rate of $1/p^2$ and since the resolution indicator works with the log of squared quantities, ψ_0 should roughly scale as $\psi_0 \sim -4 \log_{10}(p)$. In this work, $\psi_0 \equiv -(4 + 4.25 \log_{10}(p))$ and $\Delta\psi \equiv 0.5$.

4.3.2 Jump Indicator

The idea to use the uniquely DG inter-element jumps as a discontinuity indicator was first proposed by Dolejsi et al. [40] and also adopted by Krivodonova et al. [82] based on the work of Adjerdid et al. [1]. For a smooth flow solution, the magnitude of the inter-element jumps should be convergent,

$$|[[g]]| = \begin{cases} \mathcal{O}(h^{p+1}), & \text{smooth flow} \\ \mathcal{O}(1), & \text{discontinuity} \end{cases}$$

where $g = g(\mathbf{u}) : \mathbb{R}^m \rightarrow \mathbb{R}$ is a state vector component or derived quantity,

Therefore, one can easily envision an indicator that measures jumps in a state quantity

or a function to denote regions near a discontinuity. Specifically, the jump indicator is cast as,

$$J_\kappa = \frac{1}{|\partial\kappa|} \int_{\partial\kappa} \left| \frac{[[g]]}{\{g\}} \right| \cdot \hat{\mathbf{n}} \, ds \quad (4.11)$$

where jumps in pressure are chosen as the functional quantity, $g(\mathbf{u})$, to locate shocks. Additionally, for the purposes of implicit linearization, the absolute value function was substituted with a \mathcal{C}^1 -continuous approximation,

$$|x| \approx \frac{x^2}{\text{sign}(x)x + \alpha},$$

where α is an input parameter. Similar to the resolution indicator, the final scaling of the jump indicator is smoothly limited by $S_\kappa = S_\kappa(J_\kappa; \theta_S, \psi_0, \Delta\psi)$, where θ_S is the same maximum value ($\theta_S = 1$) and ψ_0 and $\Delta\psi$ are empirically determined constants, different from those of the resolution indicator. This work found that $\psi_0 \equiv -(2.25 + 3 \log_{10}(p))$ and $\Delta\psi \equiv 0.5$ were reliable quantities.

It is important to note that the resolution indicator is a function of the state in a single element. The jump indicator, however, is dependent on the state values in neighboring elements as well. In the non-smooth viscosity approach to shock capturing, using the jump indicator with an otherwise compact discretization of diffusion terms would expand the numerical stencil of the entire scheme. This is because the artificial viscosity that is applied along element edges becomes dependent on the state values in immediate and second-degree neighbor elements as well. In contrast, with the artificial viscosity equation, the jump indicator is a source function and does not spread the numerical footprint of the scheme.

4.4 Artificial Viscosity Model Comparisons

The above sections described two different artificial viscosity models: non-smooth and PDE-based. This section presents test cases designed to compare and contrast the performance of the two models.

4.4.1 Convergence Rate Accuracy

Smooth Flow

Both the resolution and jump indicators are designed to highlight under-resolved flow regions, such as those in the proximity of a discontinuity, that require the addition of artificial viscosity. However, for smooth, resolved flows, the non-linearity of the shock indicators should not flag any troubled cells so that artificial viscosity is not unnecessarily added to the discretization.

The preservation of accuracy in smooth flow was tested on the problem of 2D, inviscid, subsonic flow over a Gaussian bump at a freestream Mach number of 0.5 and zero angle of attack. An accuracy study of the global entropy norm, $\|s - s_\infty\|_2$, was performed over five grids representing four uniform grid refinements from 400–102,400 elements. The 1600-

element mesh is shown in Figure 4-3. The total temperature, total pressure and flow angle were specified at the inflow boundary and the ambient static pressure was specified at the outflow boundary. Flow tangency was enforced at the upper and lower domain boundaries, and the bump surface was approximated using cubic geometry elements.





Figure 4-3: Gaussian bump domain and mesh for smooth flow, shock indicator accuracy study (1600 elements).

The nominal grid convergence, without the use of shock capturing, is compared to the non-smooth and PDE-based artificial viscosity methods in Figure 4-4 using both shock indicators. The exact value for the entropy norm was taken from a $p = 3$ solution on a 409,600-element mesh. The nominal case achieves the optimal $\mathcal{O}(h^{p+1})$ accuracy for all values of p at every grid refinement, except for $p = 3$. With sufficient flow field resolution, both of the artificial viscosity models also achieve the same accuracy, independent of the shock indicator used. For the low-order solutions on the coarsest mesh, where the flow is not well resolved, there are small quantities of artificial viscosity added to the discretization. The shock indicators are performing as desired, flagging the under-resolved elements. However, with sufficient resolution in h and/or p , no artificial viscosity is added to the flow and the optimal accuracy is recovered.


Discontinuous Flow

In addition to their smooth flow behavior, it is also desirable for the shock indicators and the artificial viscosity models to attain analytic convergence rates for flows with discontinuities as well. For discontinuous functions, the optimal, convergence rate in the L^1 norm for an optimal L^1 polynomial approximation is $\mathcal{O}(h/p)$ [56]. This is because the L^1 -error is dominated by the discontinuity, which has zero thickness. Thus, the convergence rate in L^1 depends on how well the solution can approximate the thickness of the discontinuity, which is governed by the resolution length scale, h/p . For a least-squares polynomial approximation, as is obtained with Galerkin finite element methods, the L^1 -error should also scale with the resolution length scale of the scheme [138].


A verification of the analytic convergence rates for flows with discontinuities was carried




(a) No artificial viscosity



(b) NS, Resolution indicator



(c) PDE, Resolution indicator



(d) PDE, Jump indicator

Figure 4-4: Grid convergence rates of global entropy norm for inviscid flow over a Gaussian bump, $M_\infty = 0.5$, $\alpha = 0^\circ$ with non-smooth (NS) and PDE-based artificial viscosity models.

out in the context of the modified Burgers equation with the discontinuous forcing function described in (3.6) and (3.7). However, instead of prescribing the artificial viscosity as an explicit function in space, the non-smooth and PDE-based artificial viscosity models were used. The results in Figure 4-5 demonstrate that the optimal rate is achieved for shock flow cases. The grid convergence rates for the test case, across three uniform grid refinements, for $p = 1-4$ are $\mathcal{O}(h)$. Multiplying each line by their respective order, p , collapses all of the lines onto one, confirming the $\mathcal{O}(h/p)$ behavior. This is true for both the non-smooth and PDE-based artificial viscosity models and for both the resolution and jump indicators.

Along with the convergence rates, it is also informative to examine the behavior of the solutions near the shock, as was done in Section 3.2. Figure 4-6 depicts the $p = 6$ solution, solution gradient and viscosity distribution near the shock for both the non-smooth and PDE-based artificial viscosity models using the resolution indicator. Although both solutions are smooth, the non-smooth viscosity solution once again exhibits oscillations in the solution gradient corresponding to large jumps in viscosity. For the PDE-based viscosity solution, the inter-element DG jumps in the working variable, ϵ , create small slope discontinuities at element boundaries. However, the solution derivative is still nevertheless much better behaved than the non-smooth viscosity solution.


4.4.2 Transonic flow: NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$

The next test case is the inviscid $p = 5$ solution of a NACA 0012 airfoil with a freestream Mach number of $M_\infty = 0.8$ at an angle of attack of $\alpha = 1.25^\circ$. The inflow was specified by the total temperature, total pressure and flow angle while the outflow was specified to be the atmospheric static pressure. The airfoil surface was approximated using cubic geometry elements.

The Mach number contours and the surface pressure coefficient are shown in Figure 4-7 for the non-smooth and PDE-based viscosity solutions using the resolution indicator. Both solutions show good definition of the strong shock on the suction side of the airfoil and of the weaker shock on the pressure side. Also, by the grid overlay, notice how the shock is captured within a single element, demonstrating the sub-cell shock capturing capabilities of higher-order artificial viscosity solutions. The shock transition can occur within one element, despite the necessary addition of the viscosity across a number of elements through the shock. Additionally, when comparing the surface pressure coefficient,


$$C_p = \frac{p - p_\infty}{\frac{1}{2}\rho V_\infty^2},$$

the solutions are also quite similar with nearly identical C_p -distributions and shock-width resolution. The PDE-based artificial viscosity solution has a slightly more smeared out suction side shock, due to the larger viscosity footprint. The drag estimates differ by 1% ($c_d = 0.0225$ versus $c_d = 0.0227$).




(a) NS, Resolution switch

$$\times p =$$




(b) NS, Resolution switch
multiplied by p




(c) PDE, Resolution switch

$$\times p =$$




(d) PDE, Resolution switch
multiplied by p



(e) PDE, Jump switch

$$\times p =$$



(f) PDE, Jump switch multiplied
by p

Figure 4-5: L^1 grid convergence rates for 1D modified Burgers equation of a forcing function with discontinuity with both non-smooth (NS) and PDE-based artificial viscosity.




Figure 4-6: Comparison of non-smooth and PDE-based artificial viscosity solutions for modified Burgers equation (40 elements, $p = 6$).

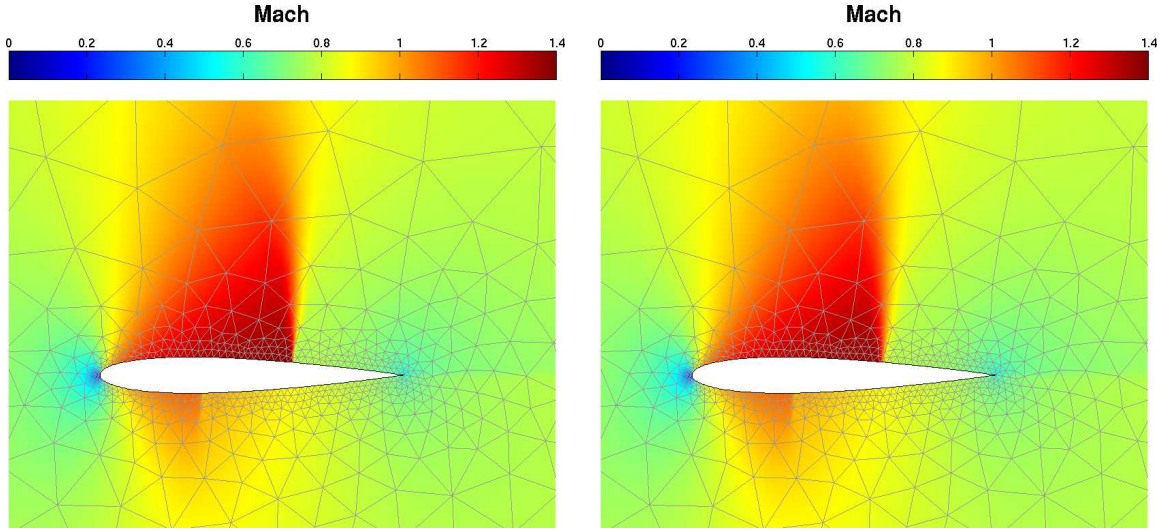
4.4.3 Supersonic flow: Half-cylinder, $M_\infty = 4$

This test case is designed to examine the effectiveness of the artificial viscosity models for stronger shocks and, in particular, focus on the behavior of the stagnation quantities behind a shock front. The case also offers the opportunity to evaluate the efforts taken to conserve total enthalpy through a shock described in Section 2.3.1. Unlike the previous examples, the benefits of the PDE-based artificial viscosity are more apparent in this application.

The solution for a 2D half-cylinder in a steady, inviscid, Mach 4 flow was solved on a sequence of three grids, shown in Figure 4-8, representing two uniform refinements. The full inflow state vector was specified at the inflow boundary and flow tangency was enforced on the cylinder surface. Since the outflow is supersonic, extrapolation boundary conditions were applied to the two outflow boundaries on either side of the cylinder.


The Mach number contours of the solution using the two artificial viscosity models and the resolution indicator are shown in the first row of Figure 4-9. The remaining rows represent the variation of total enthalpy behind the shock along the solid black line in the Mach number plots. Using the Roe flux [120] and the Laplacian artificial viscosity matrix, \mathbf{A}_ϵ in (2.9), produces significant variations in total enthalpy behind the shock, with more oscillation in the non-smooth artificial viscosity than the PDE-based method. If the artificial viscosity matrix designed for the preservation of total enthalpy, $\tilde{\mathbf{A}}_\epsilon$ in (2.10), is applied instead, the variation in total enthalpy is significantly reduced for both viscosity models (notice the change in axis scaling). Finally, when the van Leer-Hänel flux function is used [59], the variation in total enthalpy is further damped. Changing the flux function has a smaller impact on the variation in total enthalpy because the relative amount of dissipation added to the scheme by the flux function is much less than the artificial viscosity matrix. The variations of total enthalpy behind the shock are more pronounced for the non-smooth artificial viscosity model than the PDE-based approach. The higher amplitude variations are also associated with the coarsest grid for both artificial viscosity models.

Figure 4-9 demonstrates that despite the efforts to use a discretization that preserves total enthalpy through a steady state shock, small variations in total enthalpy remain. This



(a) Mach, non-smooth

(b) Mach, PDE-based



(c) Pressure coefficient

Figure 4-7: Inviscid $p = 5$ solution with resolution shock indicator of a NACA 0012 airfoil, $M_\infty = 0.8$ and $\alpha = 1.25^\circ$.




Figure 4-8: Three grids across two uniform grid refinements used for inviscid flow over a 2D half cylinder.

is certainly evident for the non-smooth viscosity solutions, but there are also numerical variations in the total enthalpy for the PDE-based artificial viscosity solutions as well that are hidden by the scale of the plots. This is due to the combination of using quadrature to approximate the integrals in the discretization and the use of ρE as the state variable in the energy equation. Section 2.3.1 described that in order to preserve total enthalpy through a steady state shock, the residual for the energy equation must be exactly equal to the residual for the continuity equation multiplied by the total enthalpy, H , in the steady state. However, by using quadrature to evaluate the integrals in the discretization, this condition is not satisfied exactly and the energy equation residual is only approximately equal to H times the continuity equation residual. Additionally, the use of conservation variables in the state vector will generally not have pointwise constant H solutions for otherwise varying polynomial states because H is a rational function of the conservative state vector. This difficulty could be overcome by using ρH as the state variable in the energy equation, although this option was not explored in this work.

As described in Section 3.4, the total pressure is also impacted by the non-smooth viscosity model. Figure 4-10 compares the variation of total pressure along the line behind the shock front for the two artificial viscosity models using the van Leer-Hänel flux function and the $\tilde{\mathbf{A}}_\epsilon$ viscosity matrix. The PDE-based artificial viscosity solution does not exhibit the oscillations that plague the non-smooth artificial viscosity solution. Instead, the total pressure varies smoothly as the shock strength changes due to curvature.

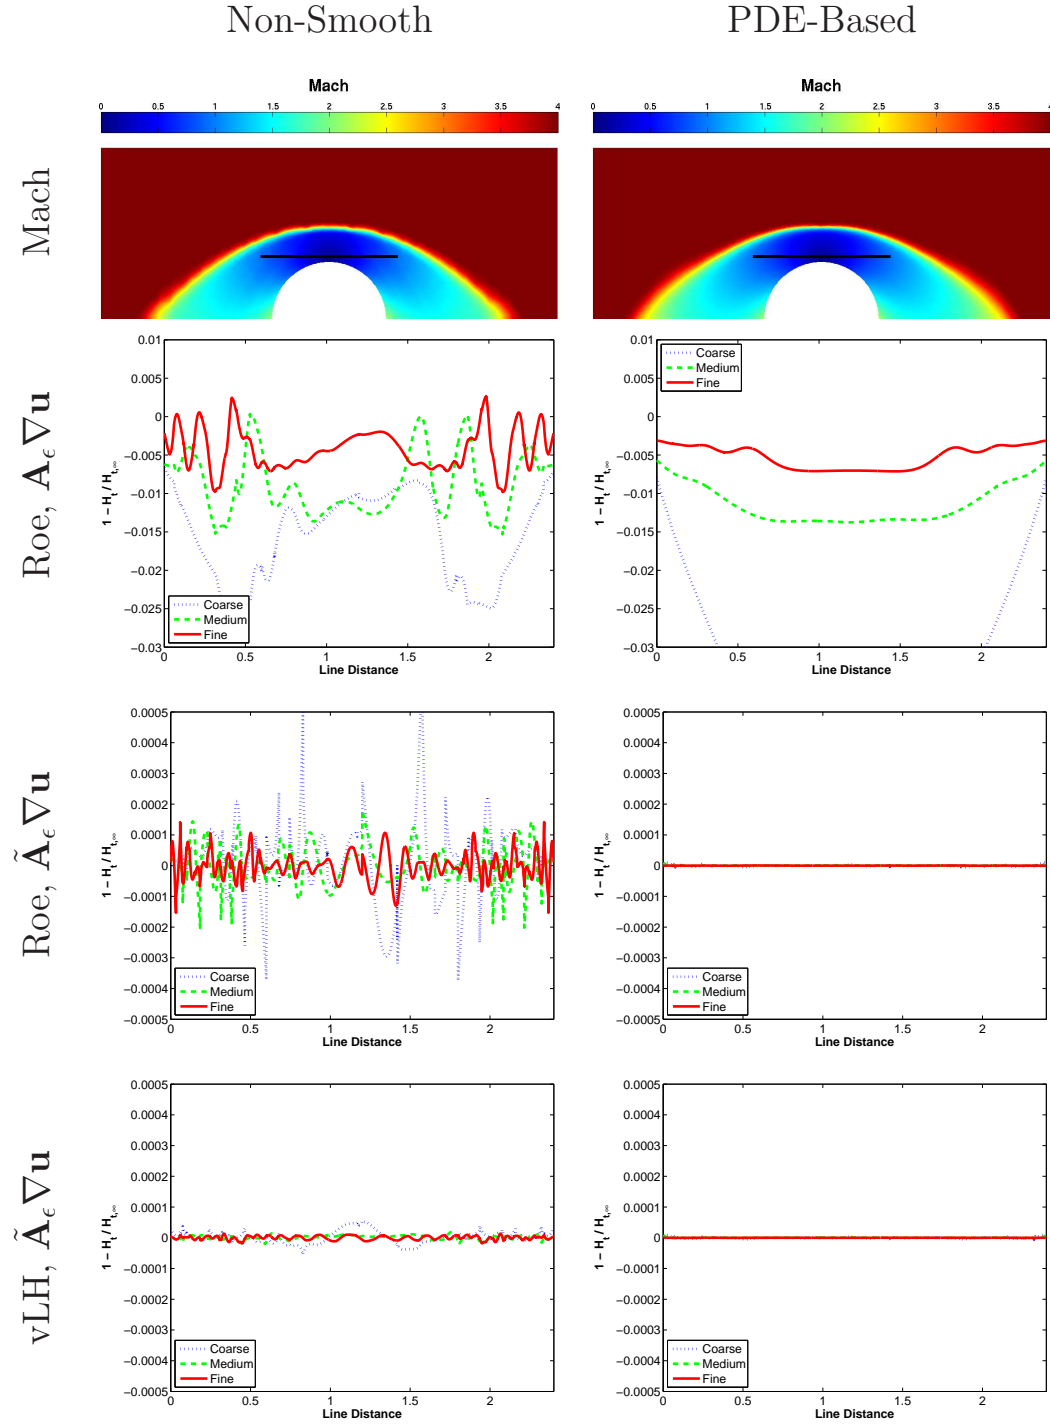


Figure 4-9: Comparison of viscosity models and shock capturing performance for a $p = 3$ solution of a 2D flow around a cylinder at Mach 4, resolution shock indicator (contour plots are shown for the intermediate mesh, vLH is van Leer-Hänel flux function).

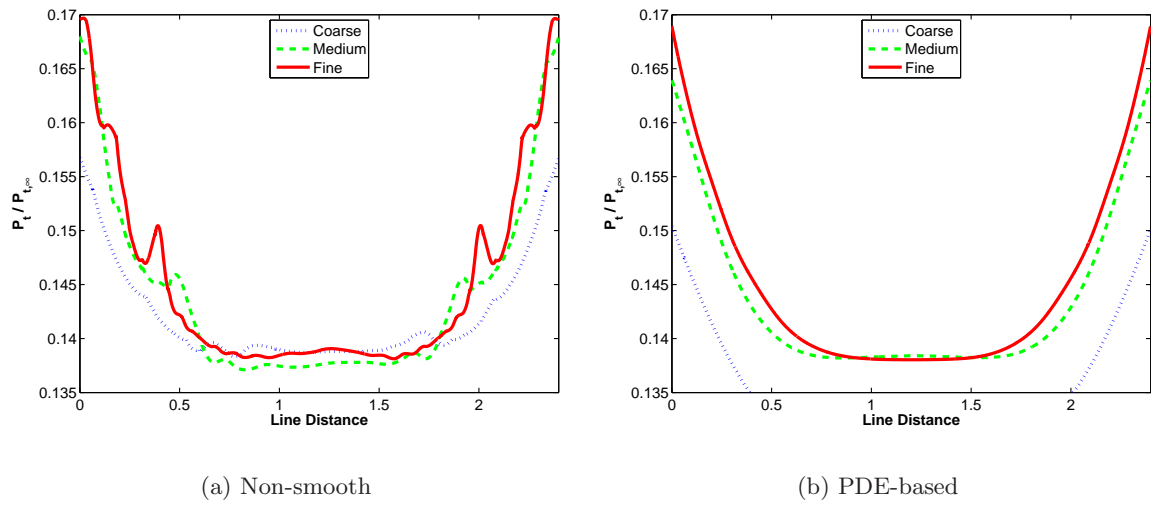


Figure 4-10: Comparison of total pressure along a measurement line behind the bow shock across two grid refinements for a $p = 3$ solution of a 2D flow around a cylinder at Mach 4 (resolution shock indicator with van Leer-Hänel flux function and $\tilde{\mathbf{A}}_{\epsilon}$ viscosity matrix).

Chapter 5

Output-Based Grid Adaptation with Shocks

Section 1.3.3 reviewed the prevalent methods for error estimation and grid adaptation. This chapter describes the approach used in this work, which is based on that of Fidkowski [43], featuring output-based error estimation and h -adaptation. First, the derivation of the error estimate and adaptation mechanics are revisited and reviewed for their applicability to discontinuous flows. The methodology is then applied to a series of example cases culminating in a 2D sonic boom model problem.

5.1 Error Estimation

In this work, the error estimation analysis and implementation is taken directly from Fidkowski [43], with minor modifications in the implementation to highlight the role of artificial viscosity. In turn, Fidkowski drew on extensive previous research by Barth and Larson [12], Becker and Rannacher [17], Giles and Süli [46], Hartmann and Houston [63], Lu [87] and Venditti and Darmofal [136, 137].

Let $\mathbf{u}_H \in \mathcal{V}_H$ be the solution to a semi-linear weighted residual,

$$\mathcal{R}_H(\mathbf{u}_H, \mathbf{v}_H) = 0, \quad \forall \mathbf{v}_H \in \mathcal{V}_H.$$

The residual function is constructed such that it accepts functions in both the discrete space, \mathcal{V}_H , and the continuous space, \mathcal{V} ; $\mathcal{R}_H : \mathcal{W}_H \times \mathcal{W}_H \rightarrow \mathbb{R}$ where $\mathcal{W}_H \equiv \mathcal{V} + \mathcal{V}_H$. Furthermore, *consistency* is assumed, in that the exact solution, $\mathbf{u} \in \mathcal{V}$, satisfies the discrete residual,

$$\mathcal{R}_H(\mathbf{u}, \mathbf{v}) = 0, \quad \forall \mathbf{v} \in \mathcal{W}_H.$$

Given a non-linear output functional, $\mathcal{J}(\mathbf{u})$, the dual problem seeks solutions, $\psi \in \mathcal{V}$, such that

$$\bar{\mathcal{R}}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{v}, \psi) = \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{v}), \quad \forall \mathbf{v} \in \mathcal{W}_H,$$

where the $\bar{(\cdot)}$ notation denotes the mean-value linearization of non-linear functions. Specif-

ically,

$$\begin{aligned}\bar{\mathcal{R}}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{v}, \mathbf{w}) &= \int_0^1 \mathcal{R}'_H[\theta \mathbf{u} + (1 - \theta) \mathbf{u}_H](\mathbf{v}, \mathbf{w}) d\theta, \\ \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{v}) &= \int_0^1 \mathcal{J}'[\theta \mathbf{u} + (1 - \theta) \mathbf{u}_H](\mathbf{v}) d\theta,\end{aligned}$$

and the primed-bracket notation denotes the Frechét derivative. Using $\mathbf{v} = \mathbf{u} - \mathbf{u}_H$ in the linearization gives,

$$\begin{aligned}\bar{\mathcal{R}}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \mathbf{w}) &= \mathcal{R}_H(\mathbf{u}, \mathbf{w}) - \mathcal{R}_H(\mathbf{u}_H, \mathbf{w}) \\ \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H) &= \mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H)\end{aligned}$$

Thus, for any $\psi_H \in \mathcal{V}_H$, the output error can be expressed as,

$$\begin{aligned}\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) &= \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H) \\ &= \bar{\mathcal{R}}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \psi) \\ &= -\mathcal{R}_H(\mathbf{u}_H, \psi) \\ &\quad + \underbrace{\mathcal{R}_H(\mathbf{u}_H, \psi_H)}_{=0, \text{by orthogonality}} \\ &= -\mathcal{R}_H(\mathbf{u}_H, \psi - \psi_H)\end{aligned}\tag{5.1}$$

By duality, this error can also be expressed in terms of the adjoint residual. Defining the adjoint residual as,

$$\bar{\mathcal{R}}_H^\psi(\mathbf{u}, \mathbf{u}_H; \mathbf{v}, \mathbf{w}) \equiv \bar{\mathcal{R}}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{v}, \mathbf{w}) - \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{v}), \quad \forall \mathbf{v}, \mathbf{w} \in \mathcal{W}_H,$$

the output error can be written as,

$$\begin{aligned}\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) &= \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H) \\ &= \bar{\mathcal{R}}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \psi_H) - \bar{\mathcal{R}}_H^\psi(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \psi_H) \\ &= \mathcal{R}_H(\mathbf{u}, \psi_H) - \mathcal{R}_H(\mathbf{u}_H, \psi_H) - \bar{\mathcal{R}}_H^\psi(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \psi_H) \\ &= -\bar{\mathcal{R}}_H^\psi(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \psi_H)\end{aligned}\tag{5.2}$$

Since the exact solutions, \mathbf{u} and ψ , are not usually known, Fidkowski employs two approximations to make the calculation of (5.1) and (5.2) practical. First, the mean-value linearization is substituted with the linearization about the discrete solution, \mathbf{u}_H , and ψ_H is therefore set to the finite element solution of the dual problem. Second, the exact solutions, \mathbf{u} and ψ , are replaced with approximations, \mathbf{u}_h and ψ_h , which exist in an enriched function space, \mathcal{V}_h . Thus, (5.1) and (5.2) are approximated by,

$$\begin{aligned}\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) &\approx -\mathcal{R}_H(\mathbf{u}_H, \psi_h - \psi_H) \\ &\approx -\mathcal{R}_H^\psi(\mathbf{u}_H; \mathbf{u}_h - \mathbf{u}_H, \psi_H)\end{aligned}$$




Figure 5-1: Comparison of exact solution approximations near a shock for the modified 1D Burgers equation, $p = 2$ solution with 50 elements.

In this work, the enriched function space, \mathcal{V}_h , is taken as the space of discontinuous polynomials of order $p + 1$, such that $\mathcal{V}_H \subset \mathcal{V}_h$. However, obtaining the exact $p + 1$ discrete solution of \mathbf{u}_h and ψ_h might be computationally prohibitive. For this reason, Fidkowski uses H^1 -patch reconstructions of \mathbf{u}_H and ψ_H in the $p + 1$ space, to obtain \mathbf{u}_h and ψ_h . This work, which includes applications with curved elements and sharp gradients near discontinuities, relies on smoothing the projections of \mathbf{u}_H and ψ_H in the enriched space. Specifically, three iterations of the element-Jacobi iterative solver are performed. Numerical tests suggest that smoothing without patch reconstruction gives better approximations to \mathbf{u}_h and ψ_h near discontinuities. For instance, consider once again the modified Burgers equation of (3.6). A $p = 2$ solution, as well as the candidate solutions in the $p + 1$ space, are shown in Figure 5-1a. Since \mathbf{u}_h is supposed to approximate the analytic, discontinuous solution, the exact $p + 1$ discrete solution performs the best in this regard. Clearly, the reconstructed solution, as a patch over neighboring elements, excessively smears the shock and is not suited for flows with discontinuities. To emphasize that the number of element-Jacobi smoothing iterations is not qualitatively important, Figure 5-1b plots the exact solution approximations obtained with five, ten, and twenty smoothing iterations- all of which lie on top of one another.

The local error indicator, in one element, is an average of the primal and dual residual expressions of the total error,

$$e_\kappa = \frac{1}{2} \left[\left| \mathcal{R}_H(\mathbf{u}_H, (\psi_h - \psi_H)|_\kappa) \right| + \left| \mathcal{R}_H^\psi(\mathbf{u}_H; (\mathbf{u}_h - \mathbf{u}_H)|_\kappa, \psi_H) \right| \right], \quad (5.3)$$

where the notation, $|_\kappa$ indicates restriction to the element, κ , and the absolute values reflect conservatism built into the error estimate. The notation, \mathcal{R}_H , reflects that the residual evaluations are done on \mathbf{u}_H and ψ_H , such that the h/p scaling of the shock switch is held

fixed regardless of the space, \mathcal{V}_h . The global output error is a summation over all elemental contributions, $e = \sum_{\kappa} e_{\kappa}$.

5.1.1 Error Estimation with Artificial Viscosity

Artificial viscosity that is non-linearly added to the governing equations creates shock thicknesses on the order of the resolution length scale of the scheme. It also fundamentally changes the governing equations of the flow field near the discontinuity. Pierce and Giles [109] used artificial viscosity for shock stabilization and separated the error into a part due to viscosity and a part due to discretization to make the error analysis more rigorous. Specifically, by looking at the discretization error for the artificially stabilized case, they regularized the problem so that the solution remains regular as $h \rightarrow 0$, and hence adequate smoothness exists to form error estimates. Similarly, the viscosity-induced error is well-behaved in the limit of vanishing viscosity. Their general analysis for quantifying the global error contributions with artificial viscosity expressed the residual of the governing equations as,

$$\mathbb{R}(\epsilon; \mathbf{u}_{\epsilon}) \equiv \mathbb{N}(\mathbf{u}_{\epsilon}) - \epsilon \mathbb{D}(\mathbf{u}_{\epsilon}) = 0$$

where \mathbb{N} is the operator of the original governing equations, \mathbb{D} is the operator for the artificial viscosity, ϵ is a constant viscosity parameter and \mathbf{u}_{ϵ} is the solution one would obtain if the viscosity were held fixed and the governing equations solved exactly. Pierce and Giles expressed the error in a functional as,

$$\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) = \underbrace{[\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_{\epsilon})]}_{\text{viscosity error}} + \underbrace{[\mathcal{J}(\mathbf{u}_{\epsilon}) - \mathcal{J}(\mathbf{u}_H)]}_{\text{discretization error}}, \quad (5.4)$$

The first term on the right-hand-side of (5.4) is the error due to the presence of the artificial viscosity and the second term is the discretization error. A Taylor expansion about \mathbf{u}_{ϵ} suggests that,

$$\begin{aligned} \mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_{\epsilon}) &= -\epsilon \frac{d\mathcal{J}(\mathbf{u}_{\epsilon})}{d\epsilon} \\ &= -\epsilon \left\langle \frac{\partial \mathcal{J}(\mathbf{u}_{\epsilon})}{\partial \mathbf{u}_{\epsilon}}, \frac{d\mathbf{u}_{\epsilon}}{d\epsilon} \right\rangle \\ &= -\epsilon \left\langle \boldsymbol{\psi}_{\epsilon}, \frac{\partial \mathbb{R}(\epsilon; \mathbf{u}_{\epsilon})}{\partial \epsilon} \right\rangle \\ &= \langle \boldsymbol{\psi}_{\epsilon}, \epsilon \mathbb{D}(\mathbf{u}_{\epsilon}) \rangle \\ &= \langle \boldsymbol{\psi}_{\epsilon}, \mathbb{R}(0; \mathbf{u}_{\epsilon}) \rangle - \langle \boldsymbol{\psi}_{\epsilon}, \mathbb{R}(\epsilon; \mathbf{u}_{\epsilon}) \rangle \end{aligned} \quad (5.5)$$

where $\langle \cdot, \cdot \rangle$ once again denotes the L^2 inner-product and $\boldsymbol{\psi}_{\epsilon}$ is the adjoint for \mathcal{J} .

If one applies the dual-weighted residual error estimate of (5.1) and (5.2) to the system of equations that includes the artificial viscosity PDE, then both the discretization error and error due to viscosity are accounted for. Since the mean-value linearization results in a bilinear operator, superposition can be used to uncover these two contributions.

Let $\mathbf{u} = [\mathbf{u}^0, \epsilon]^T$, where $\mathbf{u}^0 \in \mathbb{R}^{m-1}$ is the original state vector and ϵ is the added

state of the artificial viscosity equation. Similarly, $\mathbf{u}_H = [\mathbf{u}_H^0, \epsilon_H]^T$ and $\boldsymbol{\psi} = [\boldsymbol{\psi}^0, \psi^\epsilon]^T$. Furthermore, separate out the contributions from the primary governing equations and the artificial viscosity equations in the semi-linear residual statement,

$$\mathcal{R}_H(\mathbf{u}_H, \mathbf{v}_H) = \mathcal{R}_H^0(\mathbf{u}_H, \mathbf{v}_H^0) + \mathcal{R}_H^\epsilon(\mathbf{u}_H, v_H^\epsilon).$$

Rewriting the derivation of (5.1) with this notation gives,

$$\begin{aligned} \mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) &= \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H) \\ &= \bar{\mathcal{R}}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \boldsymbol{\psi}) \\ &= \bar{\mathcal{R}}_H^0(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \boldsymbol{\psi}^0) + \bar{\mathcal{R}}_H^\epsilon(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \boldsymbol{\psi}^\epsilon) \\ &= \bar{\mathcal{R}}_H^0 \left(\begin{bmatrix} \mathbf{u}^0 \\ \epsilon \end{bmatrix}, \begin{bmatrix} \mathbf{u}_H^0 \\ \epsilon_H \end{bmatrix}; \begin{bmatrix} \mathbf{u}^0 - \mathbf{u}_H^0 \\ \epsilon - \epsilon_H \end{bmatrix}, \boldsymbol{\psi}^0 \right) \\ &\quad + \bar{\mathcal{R}}_H^\epsilon \left(\begin{bmatrix} \mathbf{u}^0 \\ \epsilon \end{bmatrix}, \begin{bmatrix} \mathbf{u}_H^0 \\ \epsilon_H \end{bmatrix}; \begin{bmatrix} \mathbf{u}^0 - \mathbf{u}_H^0 \\ \epsilon - \epsilon_H \end{bmatrix}, \boldsymbol{\psi}^\epsilon \right) \end{aligned} \quad (5.6)$$

In the continuous limit, as $h \rightarrow 0$, the shock indicators converge to zero. Without a source term, the homogeneous solution to the artificial viscosity equation is $\epsilon(\mathbf{x}) = 0$. Thus, the exact solution vector is $\mathbf{u} = [\mathbf{u}^0, 0]^T$. Additionally, the output is only a function of \mathbf{u}^0 , $\mathcal{J} = \mathcal{J}(\mathbf{u}^0)$. With $\epsilon(\mathbf{x}) = 0$ and the output only dependent on \mathbf{u}^0 , the adjoint system reduces to that of the original governing equations and $\boldsymbol{\psi} = [\boldsymbol{\psi}^0, 0]^T$. The substitution of $\epsilon = \psi^\epsilon = 0$ into (5.6) leaves,

$$\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) = \bar{\mathcal{R}}_H^0 \left(\begin{bmatrix} \mathbf{u}^0 \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{u}_H^0 \\ \epsilon_H \end{bmatrix}; \begin{bmatrix} \mathbf{u}^0 - \mathbf{u}_H^0 \\ -\epsilon_H \end{bmatrix}, \boldsymbol{\psi}^0 \right) \quad (5.7)$$

$$= \mathcal{R}_H^0 \left(\begin{bmatrix} \mathbf{u}^0 \\ 0 \end{bmatrix}, \boldsymbol{\psi}^0 \right) - \mathcal{R}_H^0 \left(\begin{bmatrix} \mathbf{u}_H^0 \\ \epsilon_H \end{bmatrix}, \boldsymbol{\psi}^0 \right) \quad (5.8)$$

This expression for the error is identical to the derivation leading to (5.1), but with expanded vector arguments. The first row of (5.7) represents the discretization error and the second row captures the error introduced by the artificial viscosity. Furthermore, the error due to viscosity in (5.8) is clearly similar to (5.5) (for Pierce and Giles, $\epsilon = \epsilon_H$). The differences between the two expressions stems from the fact that (5.8) also includes the discretization error and in the use of \mathbf{u}^0 and $\boldsymbol{\psi}^0$ or \mathbf{u}_ϵ and $\boldsymbol{\psi}_\epsilon$. This difference would manifest itself when approximating \mathbf{u}_h and $\boldsymbol{\psi}_h$ in the computation of (5.3). If the continuous solution should be $\mathbf{u}_h \approx \mathbf{u} = [\mathbf{u}^0, 0]^T$, then \mathbf{u}_h and $\boldsymbol{\psi}_h$ should use an $h/(p+1)$ scaling of the shock switch since $\lim_{p \rightarrow \infty} \epsilon = 0$. However, if the continuous solution is $\mathbf{u}_h \approx \mathbf{u}_\epsilon = [\mathbf{u}_\epsilon, \epsilon_H]^T$, then \mathbf{u}_h and $\boldsymbol{\psi}_h$ should use an h/p scaling of the shock switch. The difference between these two approaches is depicted in Figure 5-2a. With less viscosity, the solution with an $h/(p+1)$ scaling of the switch has a sharper shock transition. Practically, however, there is little difference between the two approaches. First, solving for the exact $p+1$ solution is not computationally feasible and, as shown in Figure 5-2b, using element-Jacobi smoothing does not accentuate the differences between the viscosity scaling as dramatically. Second,

even if the exact $p + 1$ solution is obtained, there is little difference in the resulting error effectivities. The 1D example problem for the modified Burgers equation in (3.6) and the output functional in (3.11) were used to compare the exact error of a DG solution versus the estimated error via (5.3). In this way, the quality of the error estimate could be evaluated for a problem with an analytic solution and practical differences between the two scalings of the shock switch could be revealed. The effectivity was measured for $p = 1$ –4 across five uniformly refined grids, where the effectivity was defined as,

$$\text{Effectivity} = \frac{\sum_{\kappa} e_{\kappa}}{\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H)},$$


the ratio of the estimated error to the exact error. The absolute values were neglected in (5.3) so that the expression was the best possible estimate of the global error. Similarly, to obtain the best possible estimates of \mathbf{u}_h and ψ_h , the exact $p + 1$ solutions were used instead of the element-Jacobi smoothing. The results are shown in Figures 5-2c–d. For a given value of p , the error effectivities are essentially constant, with values decreasing slightly from 0.2–0.7 for increasing p . This implies that the output error estimate is converging with h in the same manner as the output itself. Furthermore, except for $p = 1$, there is little difference in the effectivities whether the $h/(p + 1)$ or h/p scaling of the shock switch is used. This supports the claim that there is no practical difference in the choice of shock switch scaling.

The error estimate used throughout this work was the expression in (5.3), the standard dual-weighted-residual error estimate of the expanded system of equations. The approximations to the continuous, exact solutions were constructed by performing three element-Jacobi smoothing iterations on both the primal and dual solutions in the enriched function space. This smoothing was performed with an $h/(p + 1)$ scaling of the shock switch. After the smoothing, the values for ϵ_h and ψ_h^{ϵ} were set to zero.


5.2 Adaptation Mechanics

As with the derivation of the error estimate, the adaptation methodology in this work is based upon that of Fidkowski, where a complete treatment can be found in [43]. The implementation has its roots in the earlier work of Castro-Díaz et al. [24], Habashi et al. [58], Peraire et al. [106], Venditti and Darmofal [136, 137] and Zienkiewicz and Zhu [146].


The 2D adaptation strategy takes a localized error estimate and uses anisotropic, h -adaptation to decrease and equidistribute the error throughout the domain. The adaptation requires a mesh size request which, in 2D, corresponds to an h_1 , h_2 , an angle, θ , and is commonly defined through a metric [24, 58]. For $p = 1$ solutions, the metric is often determined from estimates of the Hessian of a scalar field. For the Euler and Navier-Stokes equations, the Mach number is commonly used as the scalar variable. The eigenvectors of the Hessian matrix provide for the orthogonal maximum and minimum stretching directions and the ratio of eigenvalues specify the aspect ratio of the cell. However, the second derivatives of the Hessian matrix are not appropriate for anisotropy detection of high-order DG solutions. For general p , the $p + 1$ st derivative measures the interpolation error of the scalar variable.




(a) Exact $p + 1$ solutions for \mathbf{u}_h



(b) Smoothing for $p + 1$ solutions for \mathbf{u}_h



(c) Effectivities using $h/(p + 1)$ scaling for \mathbf{u}_h and ψ_h



(d) Effectivities using h/p scaling for \mathbf{u}_h and ψ_h

Figure 5-2: Additional comparison of exact solution approximations near a shock and error estimate effectivities for the modified 1D Burgers equation, $p = 2$ solution with 50 elements, with discontinuous solution.

Thus, the stretching ratios and directions are taken from the $p + 1$ derivative of the Mach number. The first principal stretching direction is the direction of the maximum $p + 1$ st derivative. The next principal stretching direction is the direction of the maximum $p + 1$ st derivative in a plane orthogonal to the first. The direction of maximum $p + 1$ st derivative is calculated in 2D by an exhaustive search of an angle range followed by bisection.

Equidistribution of the error requires a prediction of N_f , the number of elements in the adapted mesh. Let n_κ be the number of adapted-mesh elements in element, κ , with $n_\kappa < 1$ indicating coarsening. For a given element, n_κ can be expressed as,

$$n_\kappa = \prod_i \frac{h_i^c}{h_i}, \quad (5.9)$$

where h_i^c is the current element size in the i -th coordinate direction and h_i is the requested element size. If the global error is equidistributed then each adapted-mesh element is permitted to contain the error level, $e_\kappa = e_0/N_f$, and each current-mesh element is allowed an error of $e_\kappa^c = n_\kappa e_0/N_f$, where e_0 is the user-requested global error tolerance. Furthermore, changes in the element size can be combined with the local error indicator to give an *a priori* estimate of the output error as well,

$$\frac{e_\kappa}{e_\kappa^c} = \left(\frac{h_0}{h_0^c} \right)^{r_\kappa + 1},$$

where r_κ is the expected convergence rate of the error indicator. From Fidkowski, $r_\kappa = s + t - 2$ for elliptic problems and $r_\kappa = s + t - 1$ for hyperbolic problems, where $1 \leq s \leq \min(p + 1, \gamma_\kappa)$ is the convergence of the primal problem and $1 \leq t \leq \min(p + 1, \gamma_\kappa^\psi)$ is the convergence of the dual problem. Here, γ_κ and γ_κ^ψ are the regularities of the primal and dual solutions, respectively. In smooth flow cases, $\gamma_\kappa = \gamma_\kappa^\psi = p_\kappa + 1$, and Fidkowski sets $r_\kappa = 1$ near geometric singularities, such as corners or trailing edges. This work also sets $r_\kappa = 1$ for elements near discontinuities, as the error is expected to converge at $\mathcal{O}(h)$, which is confirmed through numerical testing. An element is declared to be near a discontinuity if,

$$\int_\kappa \epsilon(\mathbf{x}) d\mathbf{x} > 10 \int_\kappa \theta_L(\mathbf{u}, \mathbf{x}) d\mathbf{x} \quad (5.10)$$

where θ_L is the minimum value from (2.9). Equating the allowable error with the expected error yields,

$$n_\kappa \frac{e_0}{N_f} = e_\kappa^c \left(\frac{h_0}{h_0^c} \right)^{r_\kappa + 1}. \quad (5.11)$$

Finally, knowing that $N_f = \sum_\kappa n_\kappa$ allows for a determination of N_f .

The 2D meshing of the computational domain is done by the *Bi-dimensional Anisotropic Mesh Generator* (BAMG) [64]. BAMG allows for an input of an existing mesh and node-defined metric to produce a new mesh. Once the new mesh is created, the solution is initialized to be an L^2 projection of the solution from the previous mesh. The cases presented in this chapter use cubic geometry elements to approximate airfoil boundaries and linear elements elsewhere. Since BAMG uses linear geometry elements, at each adaptation

iteration, the higher-order nodes were inserted on a cubic spline of the airfoil geometry. The flow cases in this chapter were sufficiently benign such that this procedure did not result in any negative volume elements.

5.3 Two Dimensional Results

The above error estimation and adaptation strategy is applied to a few example problems in this section involving inviscid, viscous, transonic and supersonic flow.

5.3.1 Supersonic Flow: Compression Ramp, $M_\infty = 12$

A great deal of information regarding the behavior of the adaptation mechanics for higher-order solutions of shocked flows is communicated by considering a simple oblique shock over a wedge in inviscid flow. The case is diagrammed in Figure 5-3, where the wedge is inclined at 15° , the freestream Mach number is $M_\infty = 12$, and the shock angle is approximately $\beta \approx 19.4^\circ$. The adapted solutions are compared against uniform refinements of structured grids.




Figure 5-3: Inviscid flow over a 15° wedge, $M_\infty = 12$.

Five nested structured meshes were used as the basis of a grid convergence study for the entropy norm in the domain, $\|s - s_\infty\|_2$. The five meshes represent four uniform refinements from 984 elements to 251,904 elements. The first two grids are shown in Figure 5-4, as the finer meshes are not discernable when printed. Due to the simplicity of the flow field, the analytic expression for the entropy norm in the domain is known and can be evaluated numerically. The convergence rates for the DG solutions are shown in Figure 5-5a.




Figure 5-4: Nested structured meshes of a 15° wedge.

The structured grid results offer the opportunity to contrast the shock resolution quality of higher-order solutions versus grid refinement of low-order solutions. For instance, consider

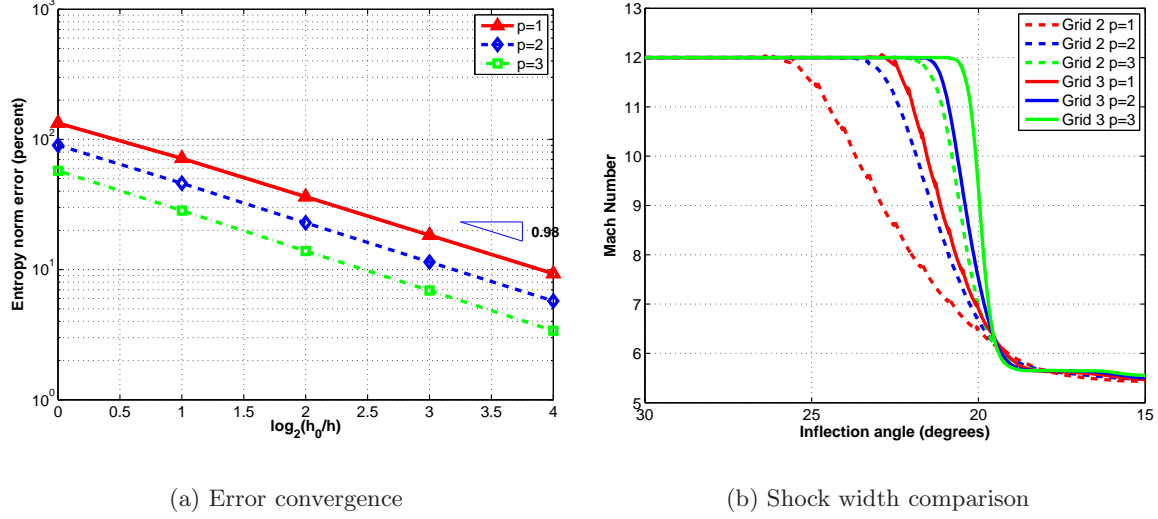


Figure 5-5: Entropy norm error and shock width convergence for inviscid flow over a 15° wedge, $M_\infty = 12$.

a second-order accurate ($p = 1$) solution on a given 2D mesh. The degree of freedom (DOF) increase for a uniform grid refinement is 400%. The DOF increase for a $p = 2$ solution on the same mesh is 200%, and for $p = 3$ is 333%. Despite the lower DOF penalty, the h/p scaling of the shock resolution suggests that a $p = 2$ solution would be roughly equivalent to the $p = 1$ solution on the finer mesh and that $p = 3$ should yield the sharpest shock jump. A sampling of the Mach number through the shock on Grid 2 and Grid 3, depicted in Figure 5-5b, corroborates the expected behavior. In this plot, the Mach number is sampled along a line in the y -direction extending from the wall half-way along the ramp. The $p = 2$ solution on Grid 2 and the $p = 1$ solution on Grid 3 have similar shock widths. The shock width of the $p = 3$ solution on Grid 2 is also approximately the shock width of the $p = 2$ solution on Grid 3.

Section 4.4.2 presented an inviscid, transonic flow case that demonstrated the sub-cell shock capturing capabilities of a $p = 5$ DG solution. Figure 5-6 depicts a zoom of the Mach number through the shock of the $p = 3$ solution on Grid 1 for the inviscid, supersonic ramp problem. Even at the high freestream Mach number of $M_\infty = 12$, the shock is resolved within 2–3 elements. Cubic interpolation is most likely not sufficient to resolve the shock within a single element.

Although higher-order solutions offer better shock resolution than grid refinements of low-order solutions, the global error is still $\mathcal{O}(h)$. To meet strict engineering error tolerances, the cell size in the shock must be reduced. This can be accomplished through uniform grid refinements. In fact, the error in the $p = 3$ solution on Grid 4, the finest structured mesh, is less than 5%. However, h -adaptation can arrive at similar error tolerances with a more economical use of computational resources. The adaptation framework was therefore directed to minimize the error in the global entropy norm, $\|s - s_\infty\|_2$, for $p = 1$ –3 until the error was below 1% of the functional value. The Grid 0 solutions served as the starting




Figure 5-6: Zoom of Mach number contours at the shock overlaid with the mesh for inviscid flow over a 15° wedge, $M_\infty = 12$ ($p = 3$ solution on Grid 1).

point for the adaptation sequences.

The final adapted meshes are shown in Figure 5-7. As expected, the grid refinement is concentrated along the shock. Figures 5-8b–c also show the convergence of the functional and the error envelope, the range of functional values obtain by considering the output plus/minus the estimated error. The error envelope, even at the earliest adaptation iterations, always encompasses the exact solution value. Figure 5-8a plots the error versus DOF for both the adapted and structured grids. From the results, it is clear that the adaptation converges towards the exact answer more efficiently than the nested structured grids. Furthermore, amongst the adapted solutions, there is a computational benefit for using higher-order ($p > 1$) interpolations, even though the flow field is constant aside from the discontinuity. The computational efficiency gain drops off though for higher values of p . This is due to the shock and can be understood by relating the error to the total DOF.

For shocked flows, assume that the global error is eventually dominated by the local errors at the shock,

$$e = \mathcal{O}\left(\frac{h}{p}\right),$$

where h refers to the cell size in the shock. The total DOF, N , for isotropic refinements scales as,

$$N = \mathcal{O}\left(\frac{1}{h^d} \left[\frac{1}{d!} \prod_{i=1}^d (p+i) \right]\right),$$

where the term in brackets expresses the DOF per element as a function of p and the expression also assumes that the shock is the dominant flow feature in the domain. Relating




Figure 5-7: Final adaptation grids for inviscid flow over a 15° wedge, $M_\infty = 12$.

the error to the DOF,


$$e = \mathcal{O} \left(\frac{1}{p} N^{-\frac{1}{d}} \left[\frac{1}{d!} \prod_{i=1}^d (p+i) \right]^{1/d} \right) = \mathcal{O} \left(N^{-\frac{1}{d}} F^{\text{iso}}(d; p) \right).$$

Table 5.1 lists the value of the function, $F^{\text{iso}}(d; p)$, for various values of p in 2D and 3D. For a fixed value of the global error, $F^{\text{iso}}(d; p)$ decreases as p increases, leading to a corresponding decrease in the DOF, N . Thus, for isotropic refinement, there is a computational efficiency benefit for using higher-order polynomials. Furthermore, the convergence rate of the error is $\mathcal{O}(N^{-1/d})$ for isotropic refinements and is exactly the rate obtained in Figure 5-8.


Table 5.1: Scaling factor of global error with respect to polynomial order for shock dominated flows.

| | F^{iso} | | F^{ani} | |
|--------------|------------------|------|------------------|----------|
| | 2D | 3D | 2D | 3D |
| $p = 1$ | 1.73 | 1.59 | 3 | 4 |
| $p = 2$ | 1.22 | 1.08 | 3 | 5 |
| $p = 3$ | 1.05 | 0.90 | 3.33 | 6.67 |
| $p = 4$ | 0.97 | 0.82 | 3.75 | 8.75 |
| $p = \infty$ | 0.71 | 0.55 | ∞ | ∞ |


For anisotropic adaptation, the local errors at the shock are dependent on the element



(a) Error convergence



(b) Functional convergence



(c) Functional envelope

Figure 5-8: Error and functional convergence histories for inviscid flow over a 15° wedge, $M_\infty = 12$.

spacing normal to the shock, recognizing different measures of h in different directions,

$$e = \mathcal{O}\left(\frac{h_n}{p}\right); \quad N = \mathcal{O}\left(\frac{1}{h_n h_2 \dots h_d} \left[\frac{1}{d!} \prod_{i=1}^d (p+i)\right]\right).$$

If the refinement is largely in h_n , such that $h_2 \dots h_d$ are effectively constant, then the error can be expressed as,

$$e = \mathcal{O}\left(\frac{1}{p} (N h_2 \dots h_d)^{-1} \left[\frac{1}{d!} \prod_{i=1}^d (p+i)\right]\right) = \mathcal{O}\left((N h_2 \dots h_d)^{-1} F^{\text{ani}}(d; p)\right).$$

For anisotropic adaptation, the global error scales with $\mathcal{O}(N^{-1})$, which agrees with the results in Figure 5-8. Table 5.1 also lists the values of $F^{\text{ani}}(d; p)$ for various values of p in 2D and 3D. In this case, $F^{\text{ani}}(d; p)$ increases for $p > 2$ in 2D and $p > 1$ in 3D. Hence, the computational benefit for using higher-order polynomials with anisotropic adaptation in the shock drops off for increasing values of p . This is a theme that is repeated in all of the applications presented in this section.

The rapid convergence of the solution using adaptation is underscored by examining the shock width resolution, shown in Figure 5-9. Whereas the shock width reduction in the structured meshes is quite gradual, the adapted grids rapidly converge to a shock that is indiscernible from a true discontinuity.

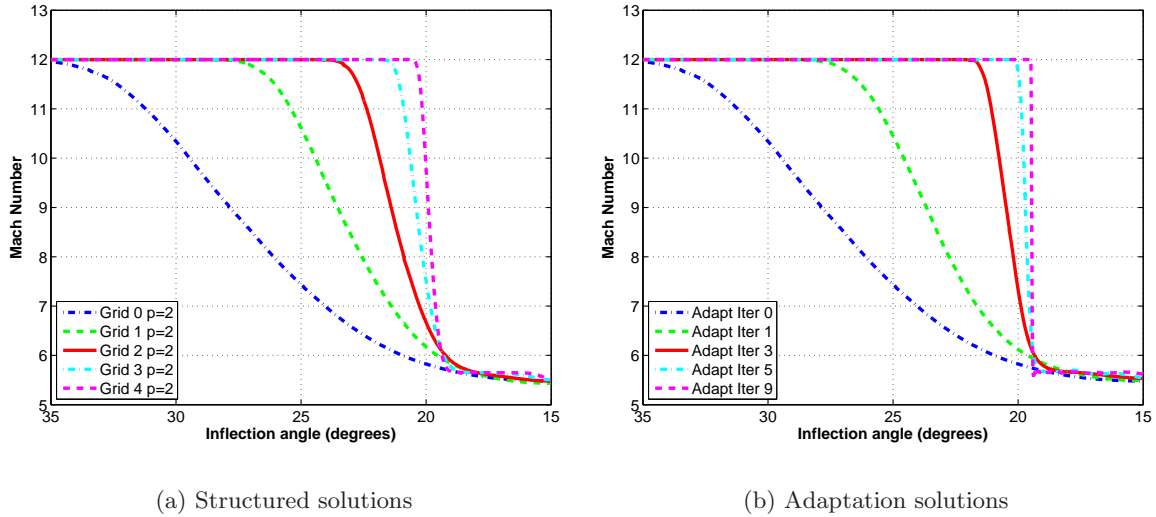


Figure 5-9: Shock width convergence of nested structured meshes and adapted grids for inviscid flow over a 15° wedge, $M_\infty = 12$.

Since the exact value of the entropy norm is known, the quality of the error estimate can be easily evaluated. The global error effectivities through the adaptation sequences, shown in Figure 5-10, are quite good. Two plots are shown. The first is the conservative error estimate of (5.3), which includes absolute values around each elemental contribution. The

second is the global error estimate without the absolute values. The conservative estimate values are generally close to 10, while without the absolute values the global errors are between 0.1–1. This is consistent with the results above in Figure 5-2.

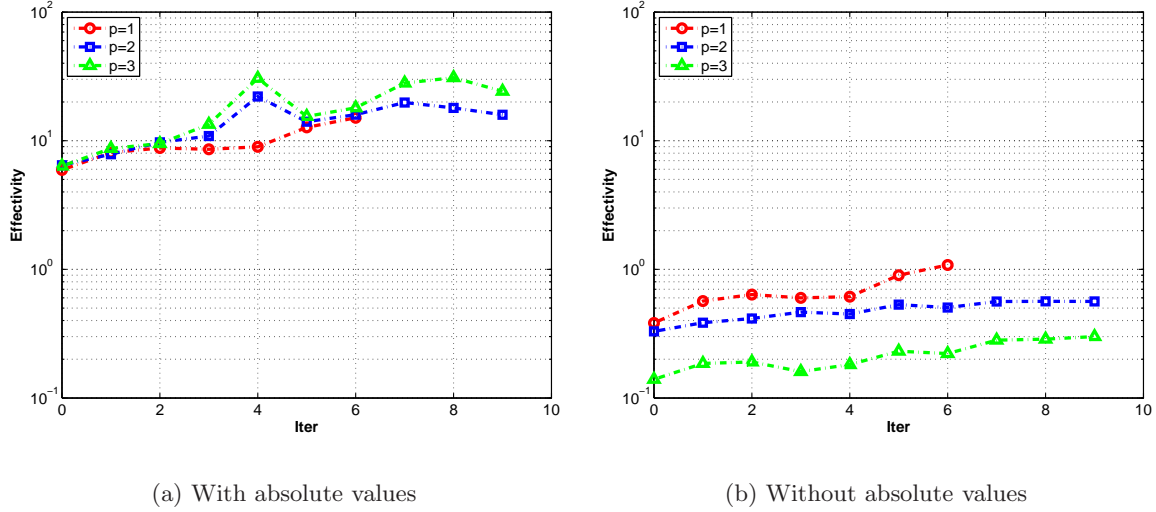


Figure 5-10: Effectivity histories for inviscid flow over a 15° wedge, $M_\infty = 12$.

Finally, Figure 5-11 depicts the region of modified regularity, defined by (5.10), due to the presence of artificial viscosity for the compression ramp case. This plot confirms that the artificial viscosity only modifies the assumed regularity in the adaptation mechanics in the vicinity of the shock.




Figure 5-11: Region of modified regularity due to artificial viscosity for inviscid flow over a 15° wedge, $M_\infty = 12$.

5.3.2 Transonic flow: NACA 0012, $M_\infty = 0.95$, $\alpha = 0^\circ$

This test case involves inviscid, transonic flow past a NACA 0012 at zero angle of attack and freestream Mach number of $M_\infty = 0.95$. The airfoil is modeled with cubic geometry elements. The adaptation is directed to minimize the error in drag on the airfoil for inter-

polation orders, $p = 1-3$. The grids were adapted until the estimated error in the drag was below 0.16 drag counts.


Far-field and near-field perspectives of the flow field of a well-resolved, $p = 3$ solution are shown in Figure 5-12. The flow is accelerated over the airfoil until it meets two oblique, fish-tail shocks that sit at the trailing edge. Behind these shocks, there is also a weak normal shock. This adaptation case was studied by Warren et al. [142] and later by Venditti and Darmofal [136], who also relied upon output-based grid adaptation. Venditti and Darmofal found that the adapted grids featured refinement at the leading and trailing edges, but little refinement of the oblique shocks. The sharp resolution of the shocks was not required for accurate estimation of the drag coefficient.

The initial grid for all p -values is shown in Figure 5-13a and the final adapted grids are in Figures 5-13b–d. At all solution orders, there is definite refinement of the leading and trailing edges. The lower order solutions also show refinement of the supersonic characteristics that determine the actual shock location on the airfoil, with the level of refinement clearly more pronounced for $p = 1$. These regions of refinement agree well with the contours of the adjoint solution, shown in Figure 5-12c, which relate local errors to the output. Since the resolution of the shocks is not necessary to accurately predict the drag on the airfoil, as shown by Venditti and Darmofal, the need for adaptation drops off for higher and higher values of p . In fact, the final $p = 3$ adapted grid is almost exclusively limited to the singularity regions.

All of the adapted solutions converge to the same value for drag, within the requested tolerance. The higher-order solutions arrive at an estimated error with fewer DOF for the same tolerance, as can be seen in Figure 5-14a. The $p = 1$ solution uses more than twice the DOF of the higher-order solutions. Once again, the efficiency gain between $p = 2$ and $p = 3$ is less than the move from $p = 1$. Figure 5-14 also includes plots of the convergence of the functional and error envelope. In these plots, the *exact* answer is taken from a $p = 3$ solution on the final $p = 3$ adapted mesh that was uniformly refined twice. One again, the error envelope encompasses the exact function value during the entire adaptation sequence.


Sensitivity to the Initial Mesh

The $p = 2$ and $p = 3$ solutions arrived at the requested drag precision with fewer DOF than the $p = 1$ solution. However, since all of the adaptation sequences were initialized from the same mesh, the higher-order solutions started with better flow field resolution than $p = 1$. To determine if the quality of the initial mesh has any bearing on the final adapted grid, the adaptation sequence was repeated and initialized on a coarser starting mesh. Whereas the original starting mesh contained 1836 elements, the coarser starting mesh contained 790 elements, with both grids shown in Figure 5-15. The error convergence versus DOF convergence history is also plotted in the same Figure. Although the starting points were different, Figure 5-15 shows that the adaptation sequences converge to the same grid resolution and DOF count, with nearly identical behaviors as the estimated error approaches the tolerance.



(a) Far-field

(b) Near-field



(c) Near-field

Figure 5-12: Near-field and far-field contour plots of Mach number and density adjoint for drag, inviscid transonic flow over NACA 0012, $M_\infty = 0.95$ (fine mesh, *truth* solution, $p = 3$).







Figure 5-13: Near-field view of initial and final adaptation meshes for inviscid transonic flow over NACA 0012, $M_\infty = 0.95$.



(a) Estimated error

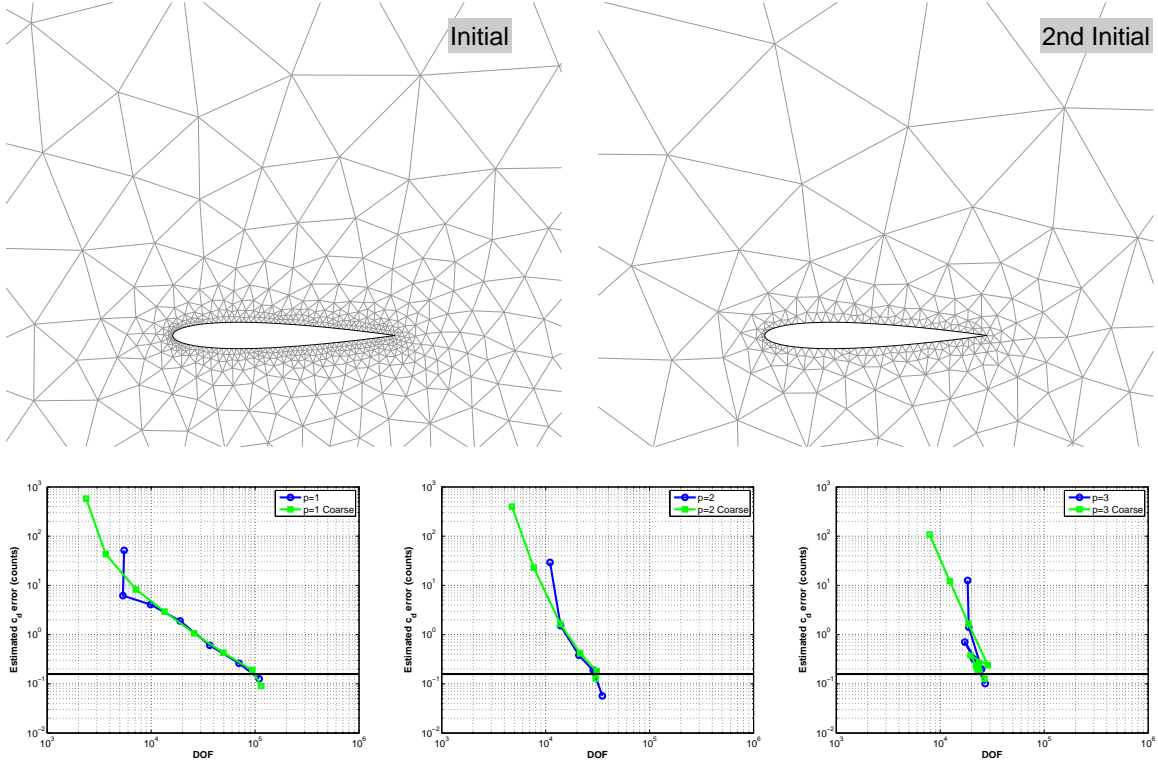


(b) c_d convergence



(c) c_d envelope

Figure 5-14: Functional and error histories of adaptation process for inviscid transonic flow over NACA 0012, $M_\infty = 0.95$.



(c) $p = 1$ estimated error

(d) $p = 2$ estimated error

(e) $p = 3$ estimated error

Figure 5-15: Sensitivity of higher-order adaptation sequence to starting mesh resolution for inviscid transonic flow over NACA 0012, $M_\infty = 0.95$.

5.3.3 Transonic flow: NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$

The next test case is also transonic flow past a NACA 0012, with a freestream Mach number of $M_\infty = 0.8$ and angle of attack of $\alpha = 1.25^\circ$ (the same case featured in Section 4.4.2). The Mach contours are reproduced here in Figure 5-16a. There is a strong shock on the suction side of the airfoil and a much weaker shock further upstream on the pressure side. The adaptation procedure for this case targeted an estimated error in lift of 0.2% c_l .




Figure 5-16: Mach and x -momentum adjoint for lift contours for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$.

The initial and final adapted grids for all values of p are shown in Figure 5-17, where the initial grid was the same as the previous test case. All adapted grids show refinement in smooth flow regions that agree with the adjoint solution contours in Figure 5-16b. This includes the leading and trailing edges, and the supersonic characteristics that determine the shock location. There is also refinement of the two shocks. However, as also determined by Venditti and Darmofal [136], resolution of the full extent of the shocks is not necessary for accurate prediction of the lift coefficient. Additionally, there are noticeably fewer elements in the smooth flow regions for higher-orders, but similar levels of refinement in the shock. The impact of these refinement features is reflected in Figure 5-18. While all of the adaptation processes converge to the same lift value within the requested tolerance, the higher-order solutions do so with fewer DOF. As before, the marginal reduction in DOF is more for the move from $p = 1$ to $p = 2$ than it is for the move from $p = 2$ to $p = 3$. Similar to the previous case, the *exact* lift value was taken from a $p = 3$ solution on the final $p = 3$ adapted mesh that was uniformly refined twice.

The region of modified solution regularity for the adaptation mechanics according to (5.10) for this test case is shown in Figure 5-19. As above, the regularity modification is restricted to the shock and does not impact the smooth flow portions of the domain.







Figure 5-17: Initial and final adaptation meshes for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$.



(a) Estimated error



(b) c_l convergence



(c) c_l envelope

Figure 5-18: Functional and error histories of adaptation process for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$.




Figure 5-19: Regions of modified regularity due to artificial viscosity for inviscid transonic flow over NACA 0012, $M_\infty = 0.8$, $\alpha = 1.25^\circ$.

5.3.4 Supersonic flow: NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$

Next, supersonic applications are considered. As described in Section 1.1.1, an adaptive method is well suited to the near-field flow analysis of a supersonic aircraft. The design of a supersonic aircraft requires accurate prediction of performance (drag) and pressure perturbations a number of body lengths away from the aircraft. These two metrics place different demands on grid resolution, which is illustrated in the next two examples.

The test case is supersonic flow over a NACA 0012 at zero angle of attack, with a freestream Mach number of $M_\infty = 2$ and Reynolds number $\text{Re} = 10^4$. The general flow field can be observed in Figure 5-20a. The airfoil is once again modeled with cubic geometry elements. The airfoil wall temperature is set to the freestream temperature.

Adaptation for Drag

The test case was first adapted for drag to a tolerance of 0.2 drag counts. The initial and final adapted grids are shown in Figure 5-21. For all grids, the bow shock ahead of the airfoil is finely resolved, but only to the extent such that the characteristics emanating from the shock impact the flow over the airfoil. As before, there is little difference in the shock refinement level for the various p -values, but there are considerably fewer elements in the boundary layer and wake for the higher-order solutions. The trailing edge shocks are also not refined to the same degree as the bow shock.

The regions of grid refinement in Figure 5-21 reflect the flow features that determine the drag on the airfoil. The sensitivity of the drag output to the flow field can be observed through the adjoint. Figure 5-20b depicts the density adjoint for airfoil drag. The cone of dependence has its apex just aft of the trailing edge and projects upstream. Within this region there are two smaller cones that relate to the bow shock position and sonic boundary near the leading edge.

The increased computational efficiency of the higher-order solutions is also apparent in Figure 5-22, which charts the adaptation convergence and the total degrees of freedom in the domain. While $p = 2$ and $p = 3$ reach the desired error level with many fewer DOF




Figure 5-20: Mach and density adjoint for drag contours of viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$.

than $p = 1$, there is little difference between the $p = 2$ and $p = 3$ DOF.

Adaptation for Far-field Pressure Integral

If the same flow conditions and geometry are instead adapted for minimal error of a pressure integral far away from the airfoil, different demands are placed on the mesh. All features of the pressure signal should be sharply resolved. Different aspects of the standard N-wave signal are more influential when considering indoor or outdoor sonic boom impact. Some important signal characteristics are peak overpressure, rise time, total impulse and shock pressure jump [123]. Additionally, some sophisticated supersonic aircraft designs intentionally deviate from the simple N-wave pressure signal. To this end, the error tolerance for the functional was set to 0.2% $\mathcal{J}(\mathbf{u})$, where the integral output for this case was,

$$\mathcal{J}(\mathbf{u}) = \int_0^L \left(\frac{p - p_\infty}{p_\infty} \right)^2 ds, \quad (5.12)$$

following the suggestion of Nemec et al. [96].

The pressure signal handed off to a sonic boom propagation code must be free of diffraction and cross-flow effects such that it can be considered a radiating source. This distance might be a few body lengths for simple geometries, but for more complex configurations, this near-field terminus might not be met for more than ten body lengths [112]. Therefore, similar to [96], the pressure integral is started twenty chord lengths above the airfoil and ten chord lengths behind it, with a total length of forty chord lengths. The exact location of the line integral can be seen in Figure 5-23a.







Figure 5-21: Initial and final meshes, for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ (adaptation for drag).



(a) Estimated error




(b) c_d convergence



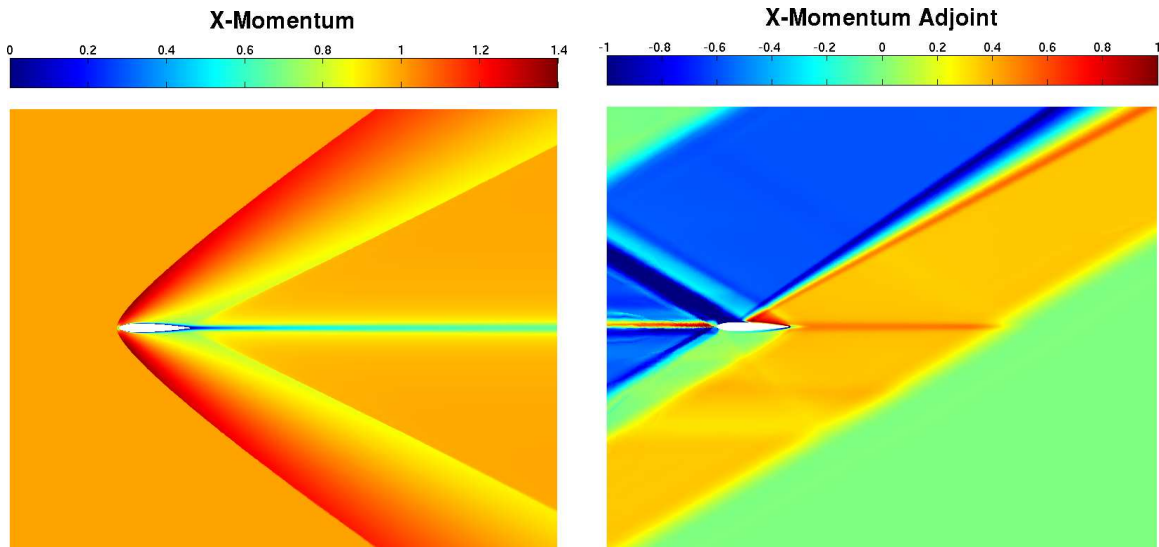
(c) c_d envelope

Figure 5-22: Functional and error histories of adaptation process for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $Re = 10^4$ (adaptation for drag).



(a) x -momentum

(b) x -momentum adjoint



(c) x -momentum near airfoil

(d) x -momentum adjoint near airfoil

Figure 5-23: Near-field and far-field contour plots of x -momentum and its adjoint for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$ (adaptation for far-field pressure).

To highlight the regions where the artificial viscosity is active, Figure 5-24 displays the regions of the flow field where the artificial viscosity is greater than the kinematic viscosity, $\epsilon(\mathbf{x}) > \nu(\mathbf{x})$. As expected, the artificial viscosity only dominates over the physical viscosity along the shocks. There is no artificial viscosity falsely added in the boundary layer near the airfoil surface or far away from the shocks.




Figure 5-24: Region of artificial viscosity greater than kinematic viscosity for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $\text{Re} = 10^4$.

The final pressure-adapted grids feature refinement in the expected areas. Specifically, the upper bow shock and trailing edge shocks are resolved as far as the pressure measurement location. This can be observed in the adapted meshes of Figure 5-25, which includes both zooms of the grid near the airfoil and measurement location. There is also some refinement of the expansion fans emanating from the curved airfoil surface and the characteristics near the measurement location. It is also interesting to note that although the pressure signal is located above the airfoil, there is still refinement below the airfoil and in the wake so that the full extent of the pressure jumps could be captured. The driving source of the adaptation in these regions can be seen by the non-zero x -momentum adjoint values in Figures 5-23b and 5-23d.

The benefits of higher-order for this case are similar to the previous adaptation examples. There is a computational efficiency benefit to using higher-order solutions, as the $p = 1$ solution meets the required tolerance with many more DOF than the higher-order solutions. However, as shown in Figure 5-26, in this case, $p = 2$ marginally outperforms $p = 3$ in computational efficiency. This is consistent with the data in Table 5.1 and the previous analysis relating the global error convergence to DOF for shocked dominated flows with anisotropic adaptation.







Figure 5-25: Initial and final meshes, for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $Re = 10^4$ (adaptation for far-field pressure).




(a) Estimated error



(b) Functional convergence



(c) Functional envelope



(d) Signal convergence, $p = 2$

Figure 5-26: Functional and error histories of adaptation process for viscous supersonic flow over NACA 0012, $M_\infty = 2$, $Re = 10^4$ (adaptation for far-field pressure).

Chapter 6

Hypersonic Applications

This chapter addresses the hypersonic flow regime. The physical processes that characterize hypersonic flow were reviewed in Section 1.1.2, along with the particular computational challenges in capturing these effects. Two application studies are presented: flow over a compression ramp and flow over a half cylinder. These applications are simple validation cases, where a comparison can be made between the work in this thesis and other computational methods as well as experimental results.

6.1 Compression Ramp, $M_\infty = 11.68$, $Re = 246,636$

The hypersonic flow over a compression ramp is a geometrically simple test case with experimental data for CFD validation studies. The geometry and flow conditions are diagrammed in Figure 6-1 and the full boundary conditions are specified in Table 6.1.




Figure 6-1: Geometry and boundary conditions for hypersonic compression ramp problem.

The problem consists of hypersonic flow over a flat plate, with length 0.442 m, that encounters a 15° compression ramp and then proceeds for another 0.269 m. The freestream Mach number is 11.68 and the Reynolds number, based on the flat plate length ahead of the ramp, is 246,636, which allows for laminar flow everywhere in the boundary layer. The gas is assumed to be thermally and calorically perfect. As shown in the contour plots of Figure 6-2, the displacement thickness of the boundary layer creates a weak oblique shock at the flat plate leading edge. A stronger oblique shock exists near the wedge corner. Additionally, the significant pressure rise behind the oblique shock at the wedge creates a strong, adverse-pressure gradient in the boundary layer. This adverse pressure gradient induces separation

Table 6.1: Freestream and boundary values for compression ramp test case.

| M_∞ | Re_L | γ | Pr | T_∞ | T_{wall} |
|------------|---------|----------|------|------------|------------|
| 11.68 | 246,636 | 1.4 | 0.71 | 64.6 K | 297.2 K |

and creates a noticeable region of recirculating flow at the wedge corner. Finally, the laminar boundary layer reattachment region corresponds to a rise in local heating which is important to predict for design purposes [65].




Figure 6-2: Contour plots of $p = 3$ solution of flow over compression ramp with $M_\infty = 11.68$, $Re = 246,636$.

Despite its simplicity, the hypersonic compression ramp is an important benchmark problem for numerical codes used for simulation of reentry environments. The interaction of shock waves with purely laminar boundary layers occurs at high altitudes, such as during space shuttle reentry [65]. Kirk [81] likens the compression ramp example to a control flap deflection of a reentry vehicle. Thus, good simulation of surface pressures is important for modeling the flap control authority and local heating must be determined for proper design of a thermal protection system.

Holden [65] first presented the experimental data for the compression ramp problem from investigations in the Calspan Corporation's 48-inch shock tunnel. The ramp model was instrumented with skin friction gauges, pressure gauges, and heat transfer gauges. The two-dimensional model was extruded in the third dimension until no three-dimensional effects

from the test section walls or model edges could be observed on the centerline measurements.

Lillard and Dries [84], and later Kirk [81], have both used this test case as a benchmark for CFD validation studies. Lillard and Dries compared results obtained using OVERFLOW, a finite-difference code, to Holden’s data. They obtained good agreement with the experimental data in terms of skin friction and heat transfer for some of the meshes they used. However, their grid converged solution deviated from the experimental data and showed a larger recirculation region than suggested by experiment. The findings of Lillard and Dries are reproduced in Figure 6-3a–b.

Kirk designed a SUPG code for reentry problems and evaluated the software on a number of test problems. For the compression ramp problem, Kirk showed very good agreement of predicted skin friction and heat transfer distributions with experimental data, as depicted in Figure 6-3c–d. It is unknown however, whether the results of Kirk were representative of a grid converged solution.


6.1.1 Structured Grid Results

The compression ramp case was computed using the PDE-based artificial viscosity model and compared against the experimental data. First, a grid convergence study was run using four nested, structured grids for $p = 1-3$, where the first three meshes are shown in Figure 6-4 (the finest mesh would appear as one solid color if printed). The coarsest mesh contained 5,184 elements and the finest mesh contained 331,776 elements. The normal wall spacing for the coarse mesh was set so that the cell Reynolds number, based on the grid spacing normal to the wall, was approximately equal to unity ($Re_{cell} \equiv \rho_w a_w h_w / \mu_w$) [53, 84].


The contour plots of the finest solution are shown above in Figure 6-2 and include depictions of the Mach number, temperature and pressure distributions. The two shocks, the first occurring at the flat plate leading edge and the second at the wedge corner, are visible in the Mach contours. Similarly, the large pressure and temperature rise behind the second shock can be observed. The separation bubble induced by the pressure rise across the shock deflects the Mach contours near the wedge corner. Figure 6-5 displays the regions of the flow field where the artificial viscosity is greater than the physical, kinematic viscosity, $\epsilon(\mathbf{x}) > \nu(\mathbf{x})$. As expected, the artificial viscosity only dominates over the physical viscosity in the vicinity of the shocks.

The evolution of the flow field resolution across refinement in both h and p is depicted in a couple of different formats. First is the comparison of Mach number contours for all p values on Grids 1 and 2 in Figure 6-6. The $p = 1$ contours for both grids do not capture the separation region at the ramp corner and the shocks are smeared over a large number of cells. The shocks are much better articulated in the higher-order solutions. The size of the separation bubble also grows with improved resolution in both h and p .


The growth of the recirculation region with improving resolution can also be observed in the plots of surface shear stress and heat transfer, shown in Figure 6-7. In these plots,




(a) Skin friction of Lillard and Dries



(b) Heat transfer of Lillard and Dries



(c) Skin friction of Kirk



(d) Heat transfer of Kirk

Figure 6-3: Comparison of computational and experimental results for the compression ramp problem obtained by Lillard and Dries [84] and Kirk [81].




Figure 6-4: Computational meshes used for grid convergence study in compression ramp problem.




Figure 6-5: Region of artificial viscosity greater than kinematic viscosity of $p = 3$ solution of flow over compression ramp with $M_\infty = 11.68$, $Re = 246,636$.

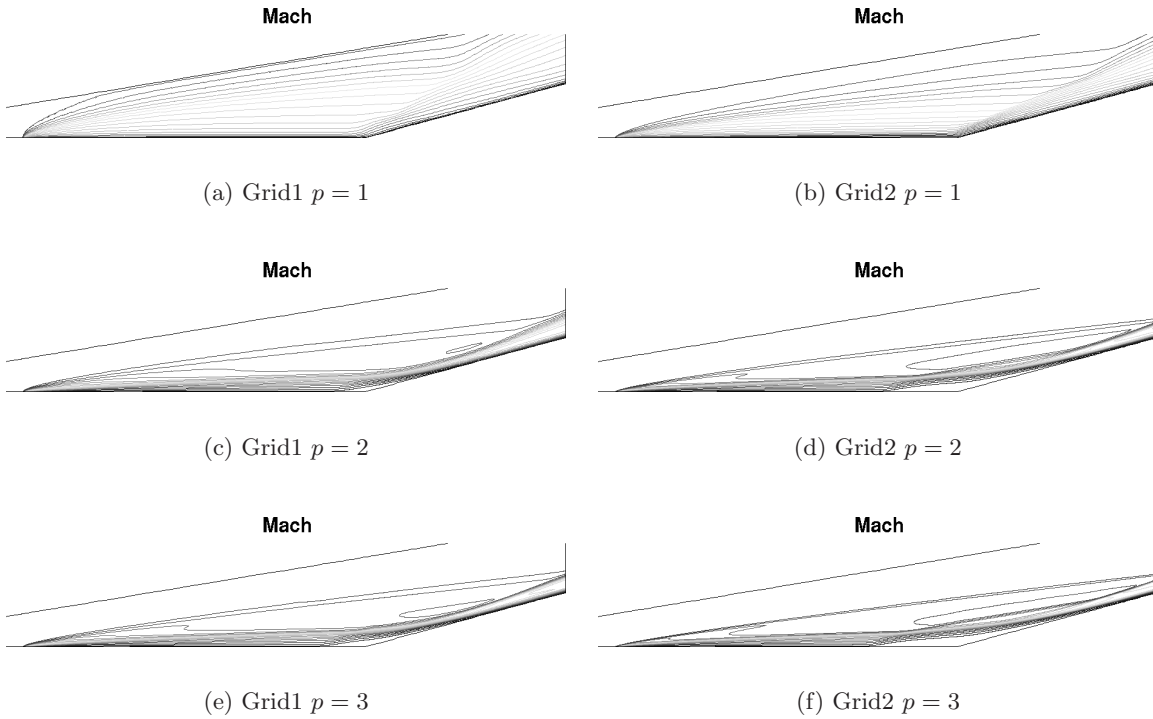


Figure 6-6: Mach contour plot convergence of $p = 1-3$ on Grids 1 and 2 for flow over compression ramp with $M_\infty = 11.68$, $Re = 246,636$.

the values along the y -axis are,

$$\text{Skin friction coefficient: } C_f = \frac{\tau_{wall}}{\frac{1}{2}\rho_\infty V_\infty^2} \quad ,$$

$$\text{Stanton number: } C_h = \frac{q_{wall}}{c_p \rho_\infty V_\infty (T_{t,\infty} - T_{wall})} \quad .$$

There is generally good agreement between the DG and experimental results up until the separation point. The skin friction plot is a good indicator of separation in the boundary layer. An inflected boundary layer profile with reverse flow yields a negative skin friction coefficient. From Figure 6-7c, not even the $p = 1$ solution on the finest mesh resolves a region of separated flow. Additionally, none of the higher-order solutions, even on the finer meshes, match the experimental data perfectly. In fact, the grid converged solutions on the finest mesh have a region of $C_f < 0$ approximately 50% larger than that of the data. This echoes the results of Lillard and Dries, who observed a grid converged solution with a larger recirculation region than that obtained by Holden. The surface heat transfer coefficient of the grid converged solution also shows some mismatch with the experimental data after the boundary layer reattachment. The reattachment location and impingement of the leading edge shock coincide with a notable increase in the heat transfer rate. The heat transfer peaks just before the exit plane of the domain is reached.

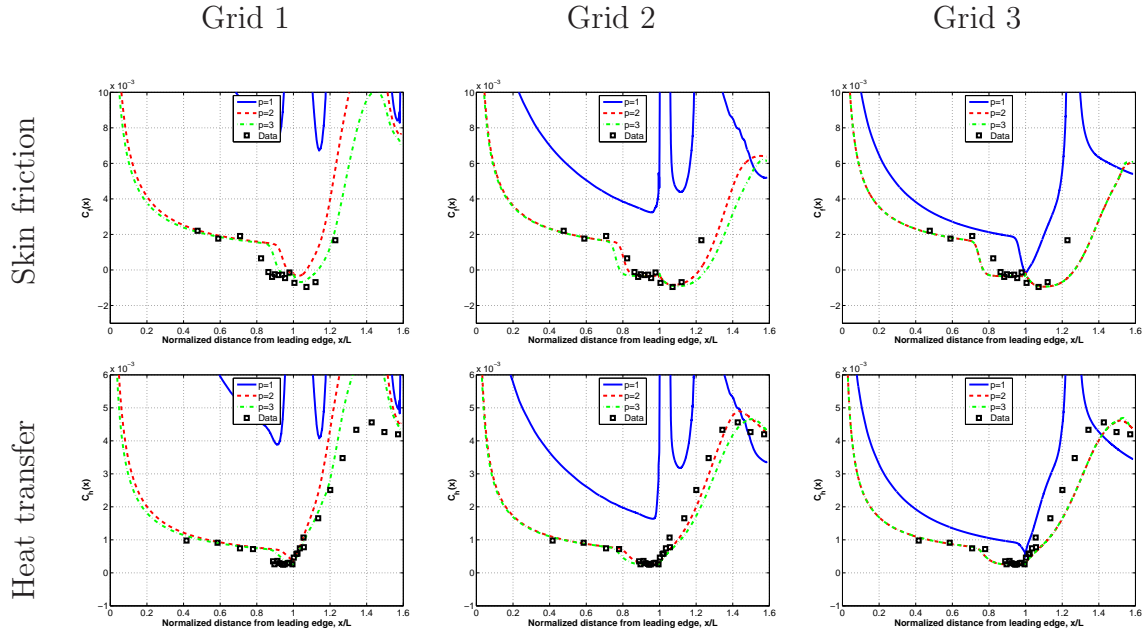


Figure 6-7: Surface plots on a given mesh from grid convergence study of flow over a compression ramp, $M_\infty = 11.68$, $\text{Re} = 246,636$.

It is difficult to ascertain with certainty the cause for the discrepancy between the experimental and computational results. The published experimental results did not include

a confidence or uncertainty level for the measurements. However, a study by Rudy et al. [121] into computational validation of laminar hypersonic compression corner flows points to the inability of capturing spanwise effects in a 2D computational simulation. Rudy et al. successfully matched experimental results using four different CFD codes for hypersonic compression corner flows with little or no separation. However, for flows with more pronounced separation, the 2D computational results over-predicted the size of the recirculation region at the compression corner. In contrast, a 3D simulation of the entire experimental test article matched the experimental data quite well. The 2D simulations did not capture the separation relief from the edge of the wedge. The results observed by Rudy et al. are quite similar to the behavior demonstrated in Figure 6-7. It is possible that this hypersonic flow case cannot be considered purely two-dimensional.

To better observe the grid convergence of the outputs, the finer solutions are overlaid on a single plot in Figure 6-8. The DOF count for the $p = 2$ solution on Grid 3 is approximately two million. The compression ramp case represents an exceedingly simple 2D geometry, but the computational expense necessary to resolve this hypersonic flow field is nevertheless exceptional. This underscores the need for automated adaptation mechanics to resolve the flow fields on more complicated 3D geometries.

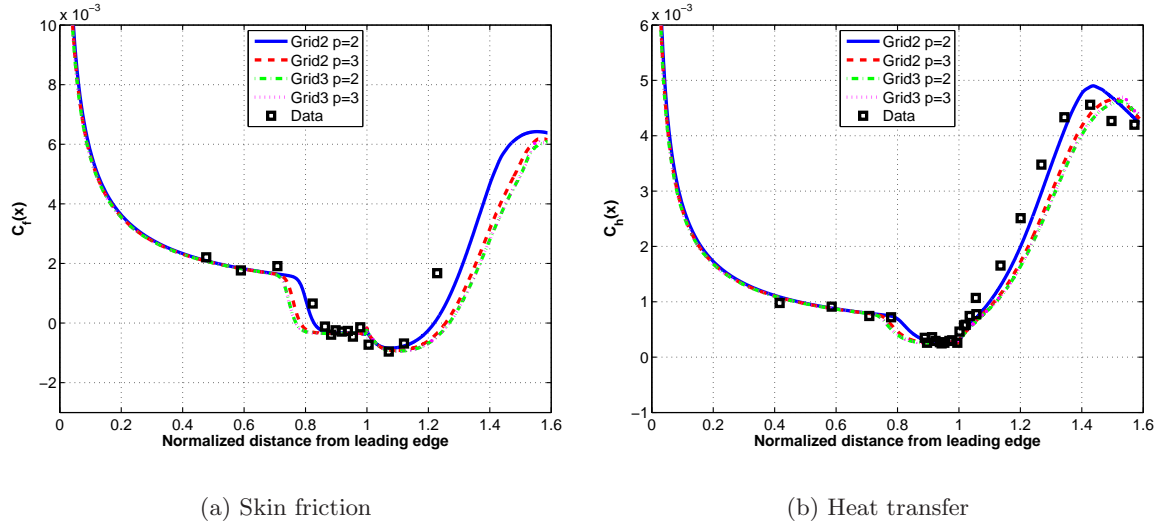


Figure 6-8: Grid converged surface plots for flow over a compression ramp, $M_\infty = 11.68$, $Re = 246,636$.

It is interesting to more closely interrogate the flow field by measuring variations of Mach number along a line extending upwards from the ramp corner. This variation of Mach number is shown in Figure 6-9. The gradual improvement in resolution from Grid 0 to Grid 3 for $p = 2$ is visible in Figure 6-9a. As the grid is refined, the shock transitions become sharper and are separated from the other flow features. Both the weak oblique shock generated at the flat plate leading edge and the stronger oblique shock near the ramp corner are visible. The passage of the measurement line through the recirculation bubble can also be observed for the higher-order solutions at the foot of the plots (separation is

not captured for the $p = 1$ solutions). The solutions for all p values on Grids 2 and 3 are compared in Figure 6-9b. This plot clearly highlights the complexities of the flow near the oblique shock that is displaced by the recirculation bubble. The subtle Mach number variations in this region are only articulated at the finest resolution levels.

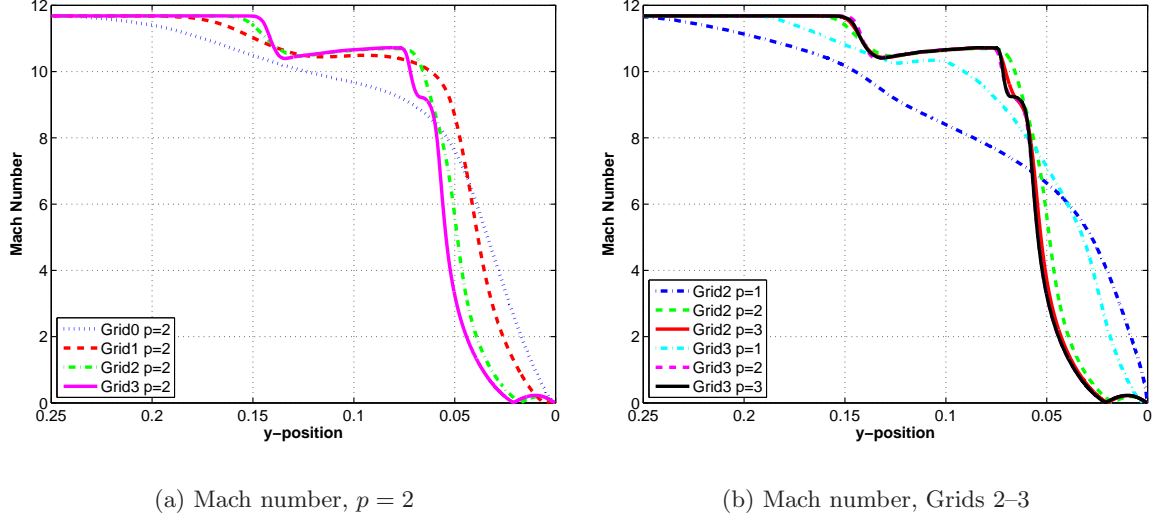


Figure 6-9: Mach number along line extending upwards from ramp corner from grid convergence study of flow over a compression ramp, $M_\infty = 11.68$, $Re = 246,636$.


6.1.2 Adaptation Results

The significant mesh size and densities necessary to approach a grid converged solution using structured meshes motivates the use of adaptation driven, unstructured meshes to arrive at the same solution. The adaptation was initialized from the solution on Grid 0 and focused on reducing the estimated error in the integrated heat flux to the surface, non-dimensionalized to be the average Stanton number on the surface,


$$Q_{wall} = \int q_{wall} ds; \quad \overline{C_h} = \frac{\frac{1}{L} \int q_{wall} ds}{c_p \rho_\infty V_\infty (T_{t,\infty} - T_{wall})} \quad .$$

The final adapted meshes are shown in Figures 6-10c-d. The adaptation clearly focused the refinement on the shocks, both the one initialized at the leading edge and at the compression corner. The final surface skin friction and heat transfer distributions are displayed in Figure 6-10, along with comparisons to the structured grid results. The adaptation results match the grid converged solution well and likewise indicate a larger recirculation region than measured by experiment. The results suggest that the adaptation framework can be applied to hypersonic flow problems and attain identical results as well designed structured meshes.


Unlike the examples in the previous chapter, the adaptation mechanics are not robust for hypersonic problems. Thus, although the final adapted solutions agreed well with the grid




(a) Skin friction



(b) Heat transfer



(c) Mesh, $p = 2$



(d) Mesh, $p = 3$

Figure 6-10: Surface plots from final adapted solution of flow over a compression ramp,
 $M_\infty = 11.68$, $Re = 246,636$.

converged solutions, the estimated error in the adaptation sequence was slow to converge and did not reach the requested error tolerance. However, despite the slow convergence of the estimated error, the adaptation sequence converged to the same heat load as the nested structured meshes using fewer DOF, as shown in Figure 6-11b. This is further supported by the surface heat transfer and Mach number measurement histories through the adaptation process, shown in Figure 6-11d–e.

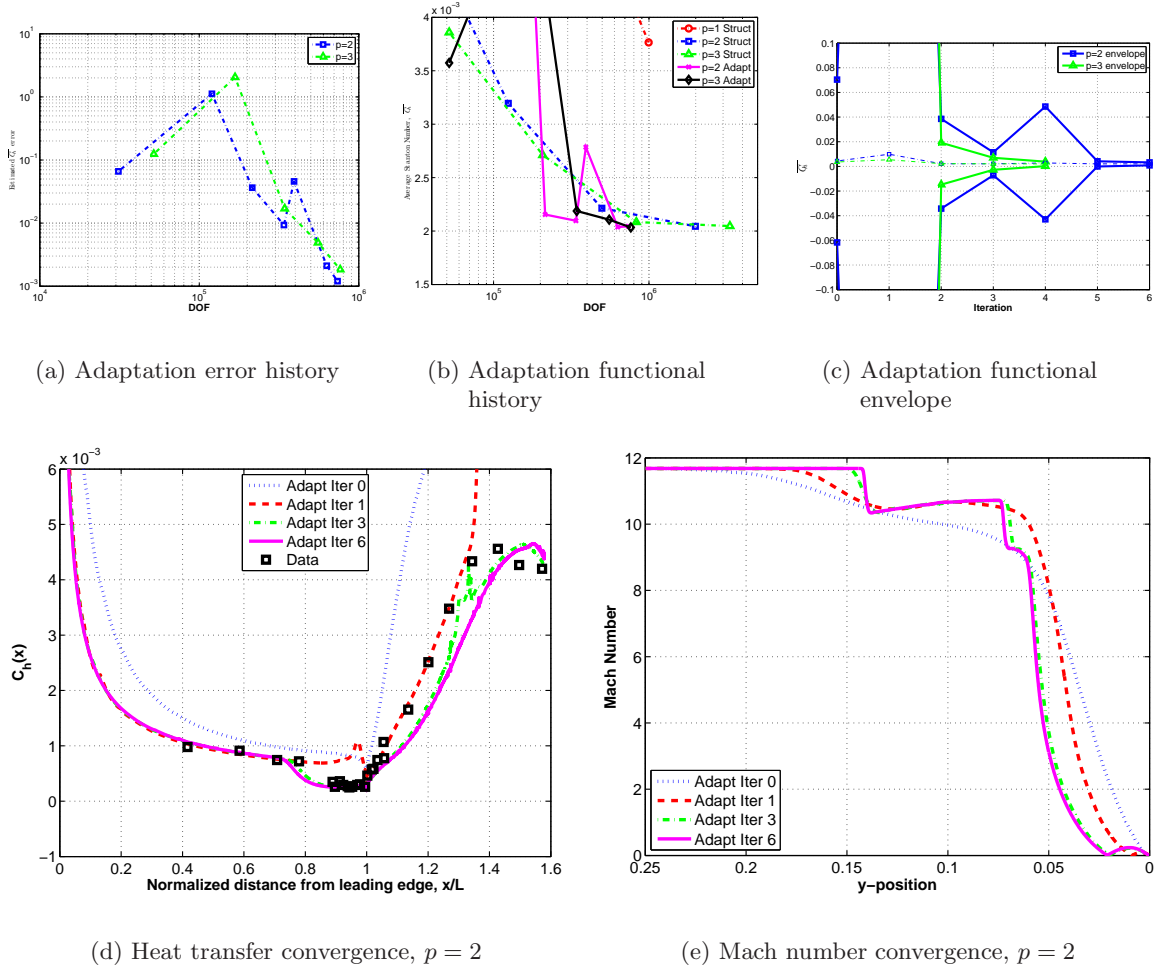


Figure 6-11: Final unstructured, adapted mesh, error history and functional history of adaptation process for flow over a compression ramp, $M_\infty = 11.68$, $Re = 246,636$.

6.2 Half Cylinder, $M_\infty = 17.605$, $Re = 376,930$

The hypersonic flow over a half cylinder is another simple test problem well suited for hypersonic CFD validation, as it highlights the difficulties in surface heat transfer prediction while using unstructured meshes. The problem is steady, features laminar flow everywhere

Table 6.2: Freestream and boundary values for half-cylinder test case.

| M_∞ | Re | γ | Pr | T_∞ | T_{wall} |
|------------|---------|----------|------|------------|------------|
| 17.605 | 376,930 | 1.4 | 0.71 | 200 K | 500 K |

in the boundary layer and the gas is modeled as thermally and calorically perfect. The complete description of the flow conditions are given in Table 6.2

6.2.1 Previous Research

Gnoffo and White [55] first presented the half-cylinder test case and compared results obtained using the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA) and the Fully Unstructured Navier-Stokes 3D (FUN3D, called HEFSS at the time of publication) codes. LAURA is a code designed for hypersonic, thermochemical non-equilibrium flows that exploits a point-implicit relaxation strategy and relies upon structured meshes [52]. FUN3D is an unstructured finite volume method that includes the ability to perform error estimation, mesh adaptation, and design optimization for fluid dynamic problems [97].

Gnoffo and White computed flow over a two-dimensional half cylinder extruded in the third dimension. For the LAURA computations, they created a structured hexahedral mesh which was adapted to align with the bow shock and included ten spanwise elements. The unstructured grid for HEFSS/FUN3D was generated directly from the structured mesh by uniformly biasing the diagonals of the cylinder surface mesh and then dividing the hexahedra into tetrahedra. The grids of Gnoffo and White are shown in Figure 6-12a–b and the surface heat transfer data at ten different spanwise locations are shown in Figure 6-12c. Unlike the LAURA results, the HEFSS/FUN3D results are asymmetrical in both the circumferential and spanwise directions. Furthermore, the variations in heat transfer near the stagnation point are as high as 20%.

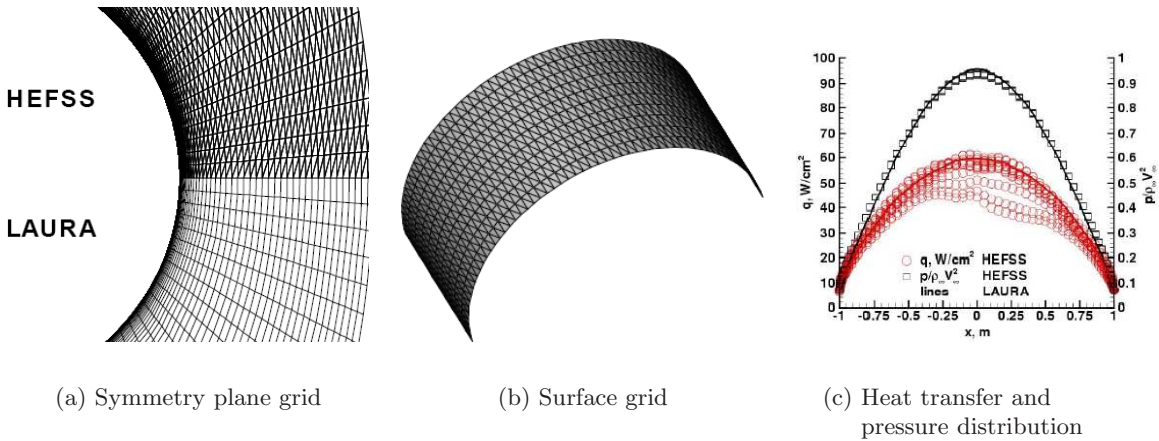


Figure 6-12: Computational mesh and results obtained by Gnoffo and White [55].

Nompelis et al. [98] replicated the 3D test case presented by Gnoffo and White [55] and further investigated the relationship between the mesh in the vicinity of the bow shock and the surface heat transfer. They constructed four different meshes, shown in Figure 6-13. The first, called the S-grid, was a hexahedral mesh with refinement near the shock. The second mesh, H1, used hexahedra in the boundary layer and shock, but tetrahedra elsewhere. The third mesh, H2, was similar to the H1-grid, but substituted tetrahedra in the shock. The final mesh, H3, used isotropic tetrahedra without additional refinement near the shock.

The surface heat transfer and pressure distributions obtained by Nompelis et al. for their four meshes are shown in Figure 6-14. Similar to the LAURA results, the structured hexahedra mesh (S) produced results with symmetrical heat transfer in both the circumferential and spanwise directions. However, slight asymmetries are introduced in the H1 mesh, when tetrahedra are used outside of the boundary layer and shock regions. For the H3 mesh, which has unstructured tetrahedra everywhere outside of the boundary layer, the asymmetries and errors in the predicted heat transfer distributions are as high as 20%.







Figure 6-13: Computational meshes used by Nompelis et al. [98].




(a) S-grid



(b) H1-grid



(c) H2-grid



(d) H3-grid

Figure 6-14: Surface heat transfer and pressure coefficient obtained by Nompelis et al. [98].

6.2.2 Discontinuous Galerkin Approach

The previous studies of the cylinder problem focused on the three dimensional setting. The DG discretization with the PDE-based artificial viscosity model and jump indicator was applied to the hypersonic cylinder problem in two and three dimensions. First, an *hp* grid convergence study was performed in 2D to confirm that the DG solutions converge to the same pressure, skin friction, and heat transfer distributions as LAURA. Next, unstructured grids were used in the external flow region to evaluate their impact on surface heating in 2D and 3D. To remain consistent with LAURA, the Roe flux was used as the numerical flux function in the DG solution as well.

6.2.3 Structured Grid Results

Four nested 2D structured meshes across three uniform refinements were used to conduct an *hp* grid convergence study. The coarsest mesh contained 4,320 elements and the finest contained 276,480 elements. The domain and first three meshes are shown in Figure 6-15. Higher-order cubic geometry nodes were inserted for every element in the domain. The full state vector was specified at the inflow boundary and a no-slip condition at the constant wall temperature was enforced on the cylinder surface. The two outflow boundaries were extrapolation for supersonic flow.

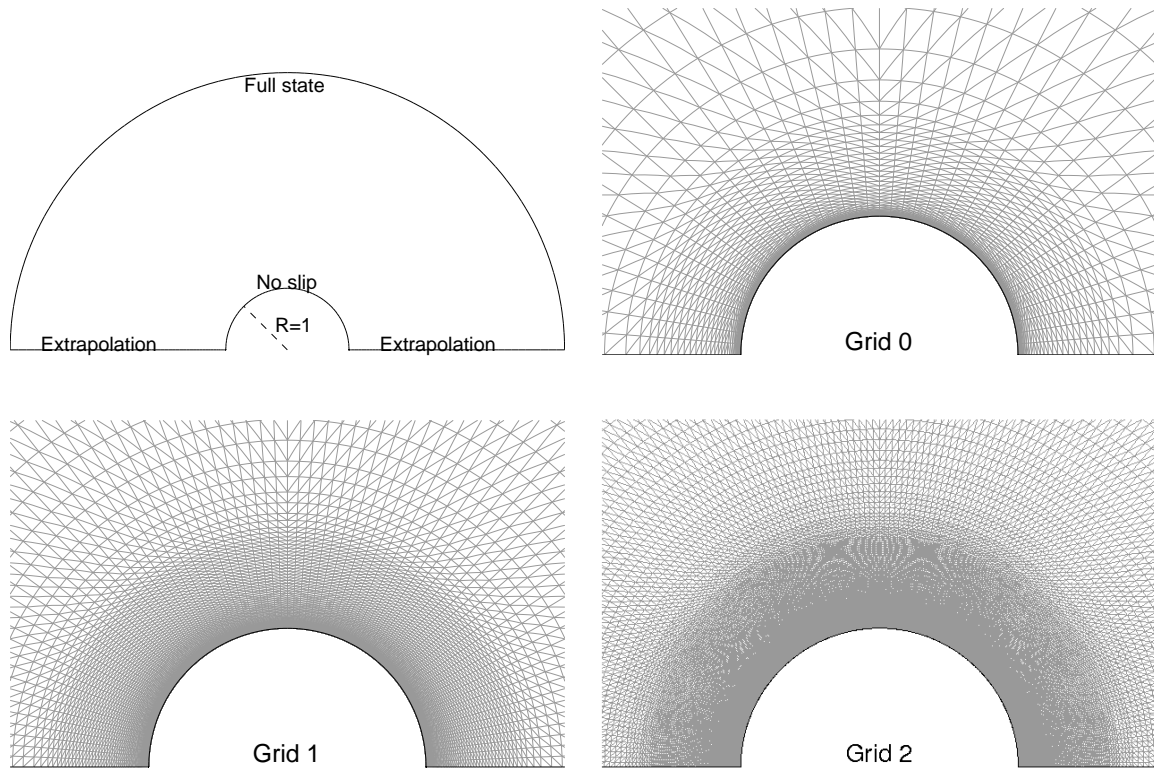


Figure 6-15: Structured grids used for 2D half cylinder grid convergence study.

The results of the grid convergence study are shown in Figure 6-16. The surface pressure distribution converges rapidly and agrees well with the LAURA results, even on the coarser solutions. The $p = 2$ and $p = 3$ solutions on the finest mesh agree well with the LAURA results for heat transfer. The coarser results tend to under-predict the peak heat transfer at the stagnation point due to the coarser resolution of the bow shock. In all of the solutions, there is a mismatch between the LAURA and DG results for surface skin friction. The grid convergence study confirms that, with sufficient shock resolution, the higher-order DG solutions with PDE-based artificial viscosity can obtain accurate surface pressure, skin friction, and heating distributions on structured meshes. The grid converged, structured mesh results will serve as a baseline of comparison for the unstructured mesh results to follow.

To emphasize that the artificial viscosity is only active near the shock, the region of artificial viscosity greater than the kinematic viscosity ($\epsilon(\mathbf{x}) > \nu(\mathbf{x})$) is shown in Figure 6-17 for the finest structured mesh solution. The highlighted region is confined to the vicinity of the bow shock and does not encroach on the boundary layer or freestream flow.

6.2.4 Unstructured Grid Results, Two Dimensions

For the 2D application, five hybrid grids were constructed, each with an identical, structured boundary layer mesh and with different unstructured meshes in the outer flow region. This focused the study on the influence of unstructured grid elements on shock resolution and on the downstream surface heat transfer. The structured boundary layer consisted of 61 nodes evenly spaced in the circumferential direction and 38 nodes in the radial direction, the position of which were borrowed from the structured LAURA grid in Figure 6-12a. The structured mesh extended to 20% of the cylinder radius in the normal direction. To build the unstructured portion of the mesh, nodes were randomly placed in the external flow region until the maximum bounding box dimension was below a prescribed threshold (0.15 was used as the tolerance). The node locations were modified slightly to eschew the creation of sliver elements. Specifically, after the addition of fifty randomly placed nodes, each node that was not located on a boundary was moved toward the center of mass of the polygon formed by the adjacent triangles. After the determination of the coordinates of the principal nodes, higher-order cubic geometry nodes were inserted for all of the elements. The near-field views of the hybrid grids are shown in Figure 6-18.

The general flow field can be observed in Figure 6-19, which includes contour plots of Mach number, temperature and pressure. Despite the coarse mesh in the vicinity of the shock and the misalignment with the bow shock trajectory, the shock is smoothly resolved. The significant temperature and pressure rise behind the bow shock can also be observed near the stagnation point. Figure 6-19d depicts contours of the added state variable, $\epsilon(\mathbf{x})$, which has units of kinematic viscosity. The contours clearly highlight the bow shock region, but are somewhat irregular as the PDE adjusts for changes in cell size and orientation along the shock.

Figure 6-20 depicts surface quantities from the five different hybrid meshes and compares them to the LAURA results. In general, the DG surface quantities agree well with the LAURA results. For the surface pressure plot, there is excellent agreement between the

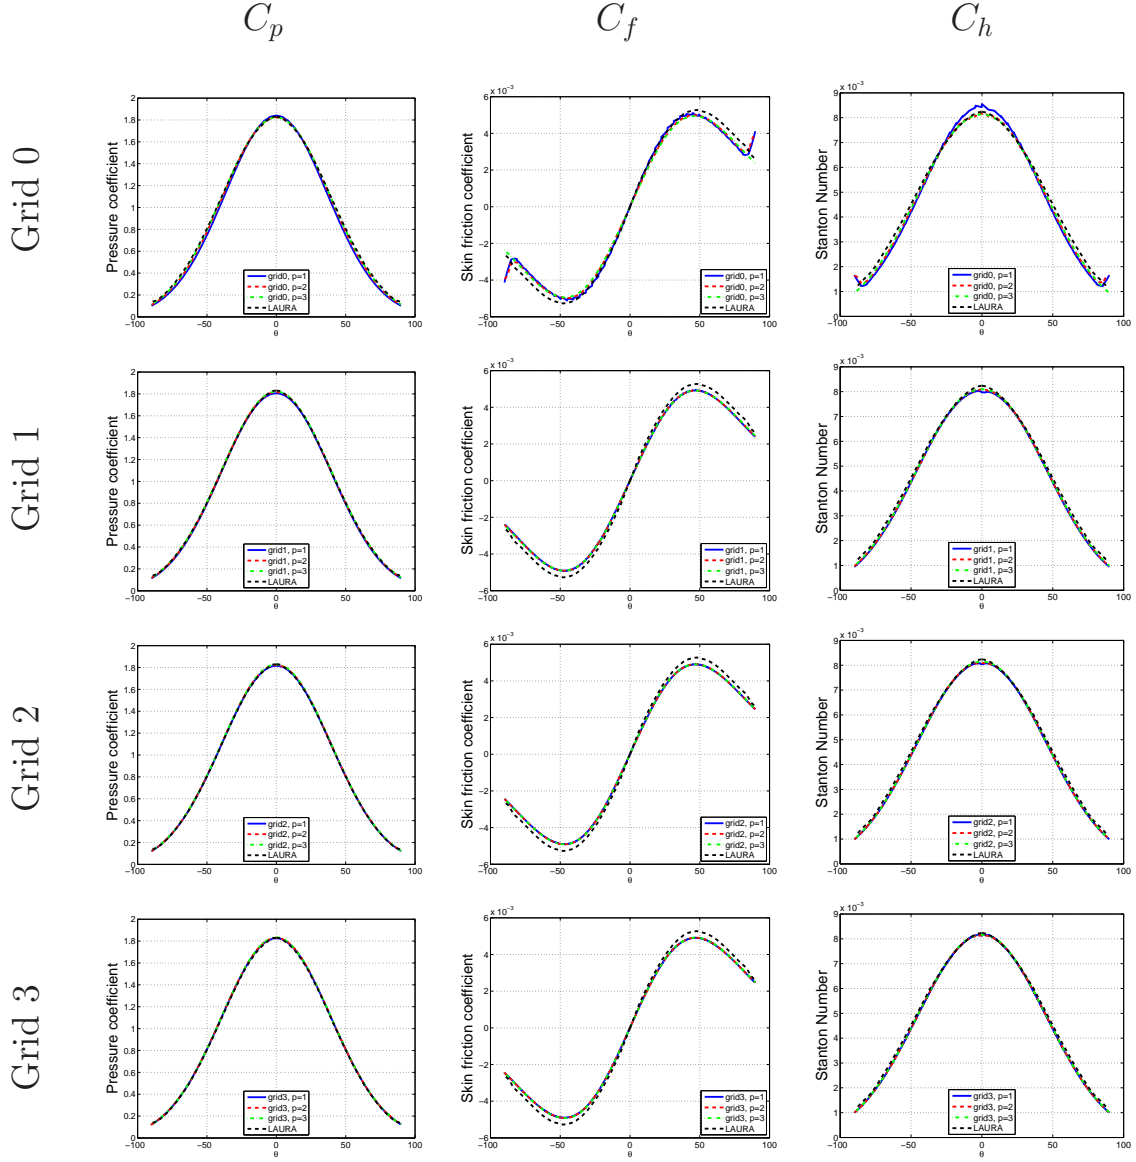


Figure 6-16: Surface plots on a given mesh from grid convergence study of flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$.




Figure 6-17: Region of artificial viscosity greater than kinematic viscosity for flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$ (taken from $p = 3$ solution on finest structured mesh).




Figure 6-18: Structured-unstructured hybrid grids used for 2D half cylinder test case.

DG and LAURA results. The only discrepancy is the pressure at the stagnation point of the cylinder, where the DG results under-predict the LAURA results by about 2%. This is perhaps due to the coarser mesh in the vicinity of the shock for the DG results. There is mismatch in the skin friction coefficient results consistent with the results from the above grid convergence study. For the Stanton number, or heat transfer coefficient, which is more sensitive to errors in shock resolution, the DG results also lie nearly on top of the LAURA results. There is a slight mismatch of the predicted peak heat transfer value and minor variation in the circumferential direction between the five different DG meshes. However, the variations are on the order of 1–2% and are far from the errors on the order of 20% reported by Gnoffo and White and Nompelis et al.. This suggests that the higher-order solution coupled with the higher-order artificial viscosity distribution has significantly mitigated the sensitivity of the shock capturing to the grid cell size and orientation along the shock.

Results demonstrating the ability of higher-order DG solutions to resolve a shock within

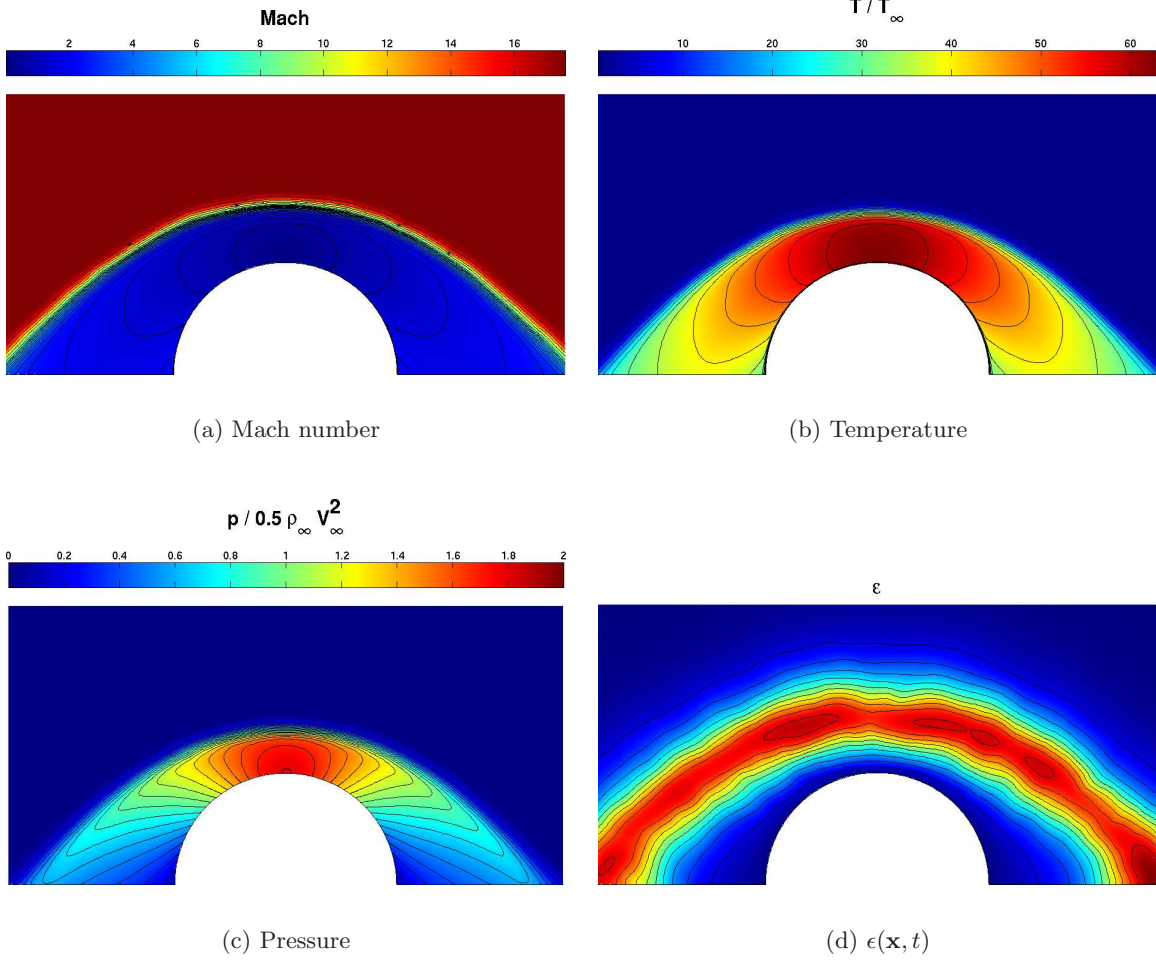


Figure 6-19: Contour plots of $p = 3$ solution of flow over 2D half cylinder with $M_\infty = 17.605$, $\text{Re} = 376,930$ (taken from Hybrid Grid 0).

a few number of elements were displayed in Sections 4.4.2 and 5.3.1. The hypersonic half-cylinder problem represents the strongest shock test case presented in this thesis. A zoom of the Mach number contours near the bow shock ahead of the stagnation point, and overlayed with the mesh, is shown in Figure 6-21. The results are taken from $p = 3$ solutions on both the structured Grid 0 and Hybrid Grid 0. For both grids the shock transition is captured in roughly 2–3 elements.

6.2.5 Unstructured Grid Results, Three Dimensions

The 3D mesh was also a hybrid grid, with a structured boundary layer mesh and an unstructured mesh in the external flow. The same structured node x - y coordinates as used in the 2D boundary layer mesh were used at six different, evenly spaced z -coordinate locations from $z \in [0, 0.5]$. This extruded the cylinder in the third dimension across five elements. The hexahedra created by the replicated node coordinates in the boundary layer were then

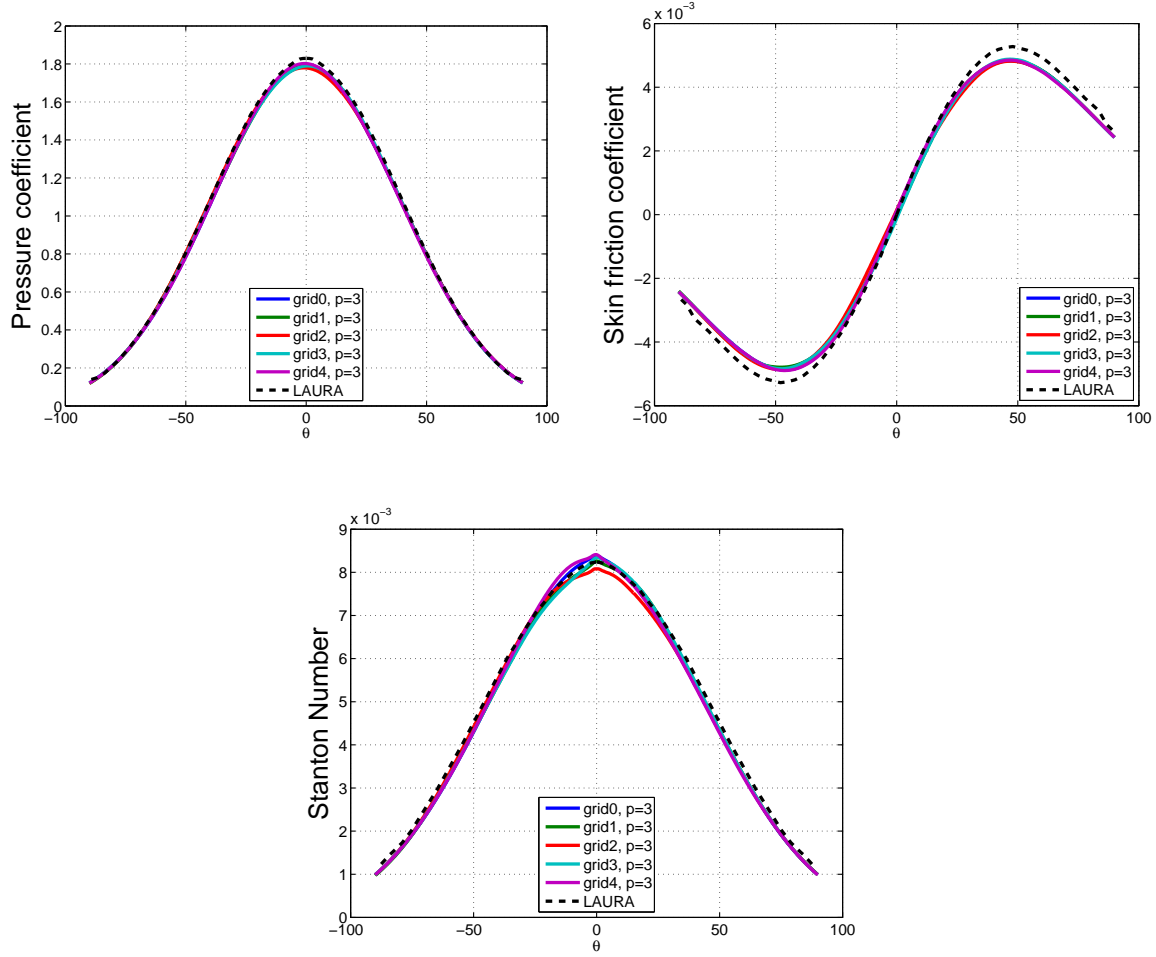


Figure 6-20: Cylinder surface plots of $p = 3$ solution of flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$ (all 5 grids shown).

split into tetrahedra. The remaining volume was filled using isotropic, unstructured tetrahedra created by *Tetgen* [126]. Finally, higher-order cubic geometry nodes were inserted for each element in the grid. Perspectives of the 3D cylinder mesh are shown in Figure 6-22. The boundary conditions are the same as the 2D problem in Figure 6-18a, with flow tangency applied to the two boundaries in the spanwise direction.

Contour plots of the 3D solution are shown in Figure 6-23, including pressure, temperature and Mach number. In these plots, half of the contours are taken from the plane $z = 0.0$ and the other are taken from the plane $z = 0.5$ to highlight any spanwise variation that might exist. As can be seen in the plots, the contours are well aligned and there is no significant spanwise variation.

The surface quantities of interest are plotted at six different spanwise locations in Figure 6-24. The results are consistent with the 2D results above, and there is generally good agreement with the LAURA results. As above, all of the surface pressure lines lie atop one another. However, there is a slight under-prediction of the pressure coefficient near the

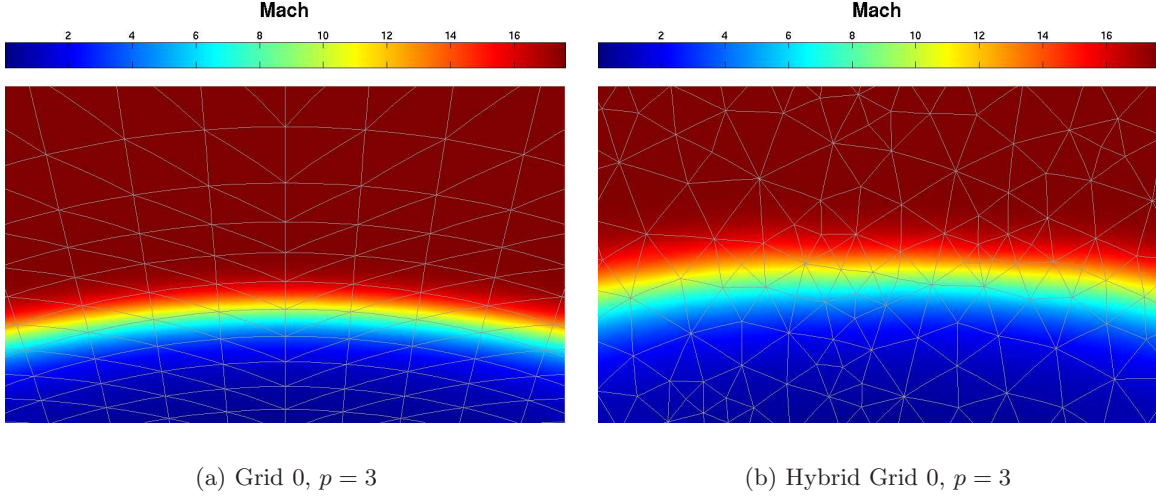


Figure 6-21: Mach contours at the bow shock of flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$ ($p = 3$ solution).

stagnation region and the mismatch in the skin friction distribution is consistent with the trend observed in Figure 6-20b. There is very little spanwise or circumferential variation in the heat transfer distribution as well, and all spanwise locations agree quite well with the LAURA predictions, despite the fully unstructured mesh in the external flow. Comparing the DG results in Figure 6-24c with the results of Gnoffo and White in Figure 6-12c or those of Nompelis et al. in Figure 6-14d clearly demonstrates the benefits of the higher-order solution and PDE-based artificial viscosity in reducing the errors introduced by the unstructured grid in the vicinity of the shock.

6.2.6 2D Adaptation

The cylinder test case was also solved using the output-based adaptation framework. Robustness deficiencies of the adaptation framework for hypersonic flows limited the solutions obtained to only $p = 2$. The flow was initialized with a structured mesh and the adaptation minimized the estimated error in the integrated heat flux to the cylinder, non-dimensionalized to be the average Stanton number on the surface,

$$Q_{wall} = \int q_{wall} ds; \quad \overline{C_h} = \frac{\frac{1}{\pi R} \int q_{wall} ds}{c_p \rho_\infty V_\infty (T_{t,\infty} - T_{wall})}.$$

The mesh generator, BAMG, uses only linear elements. For the NACA 0012 examples in the previous Chapter, the Reynolds number was small enough and the curvature of the geometry was mild enough such that using higher-order geometry interpolation only for the boundary elements did not result in negative volumes. However, for the cylinder problem, the thin boundary layer and curvature of the geometry necessitated using higher-order geometry elements everywhere in the domain. To accommodate BAMG's use of linear elements, the mesh adaptation and generation was performed in a mapped space using the

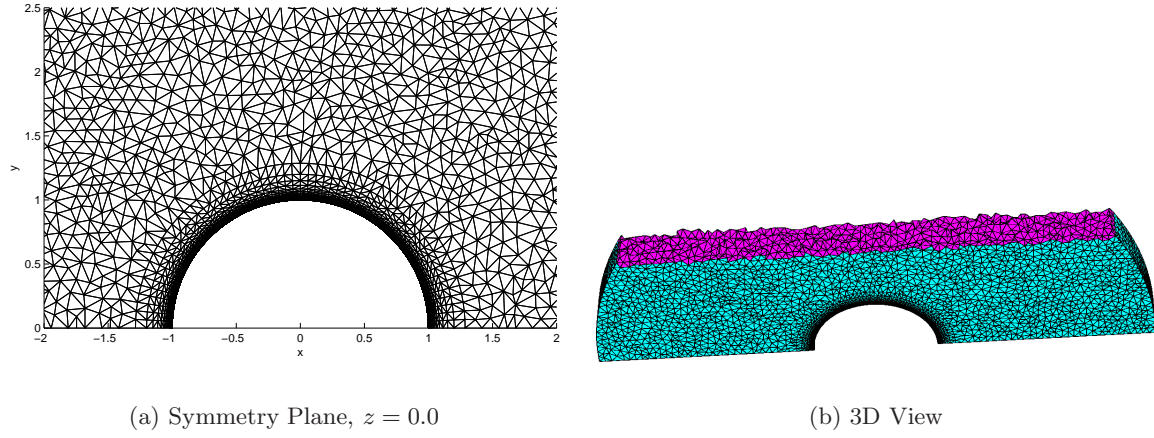


Figure 6-22: Structured-unstructured hybrid grid used for 3D extruded half cylinder test case.

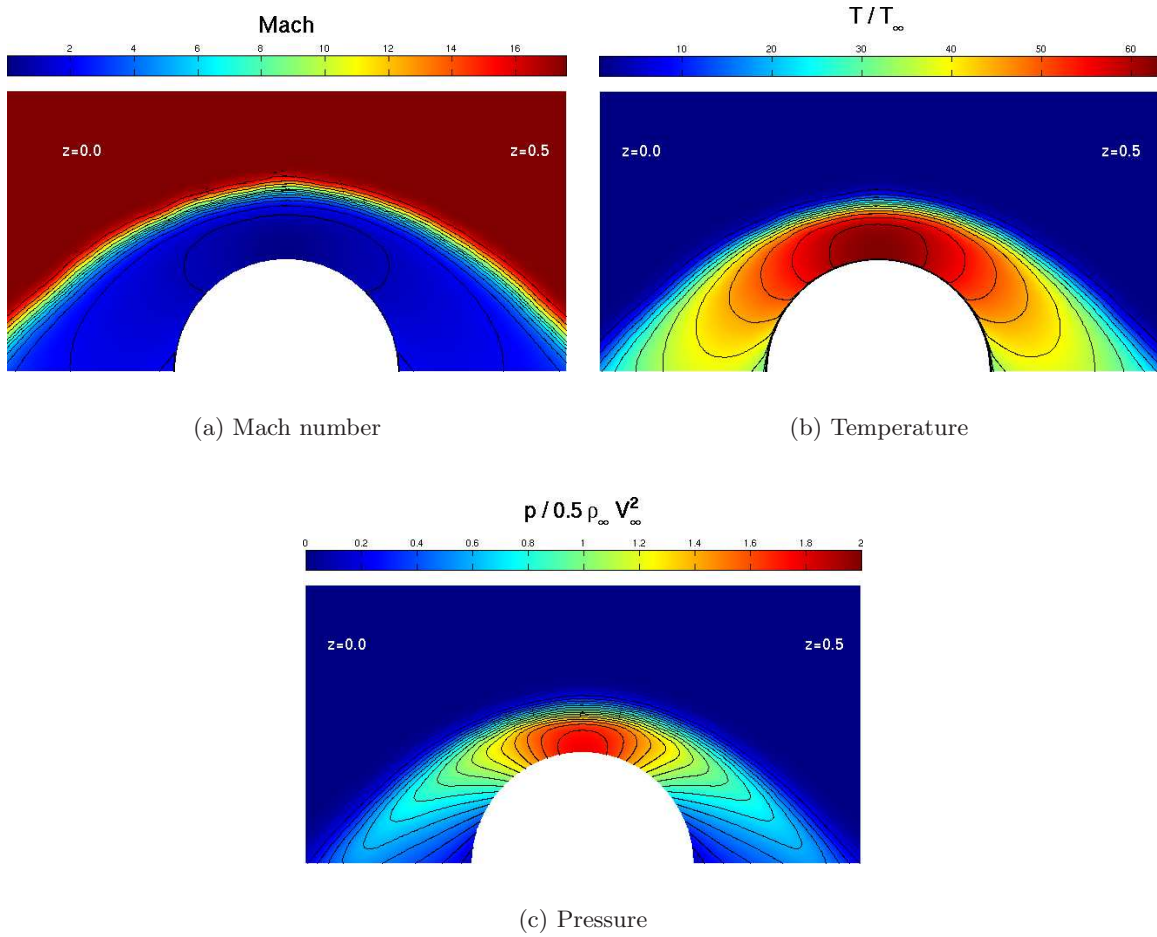


Figure 6-23: Contour plots of $p = 3$ solution of flow over 3D extruded half cylinder with $M_\infty = 17.605$, $Re = 376,930$.

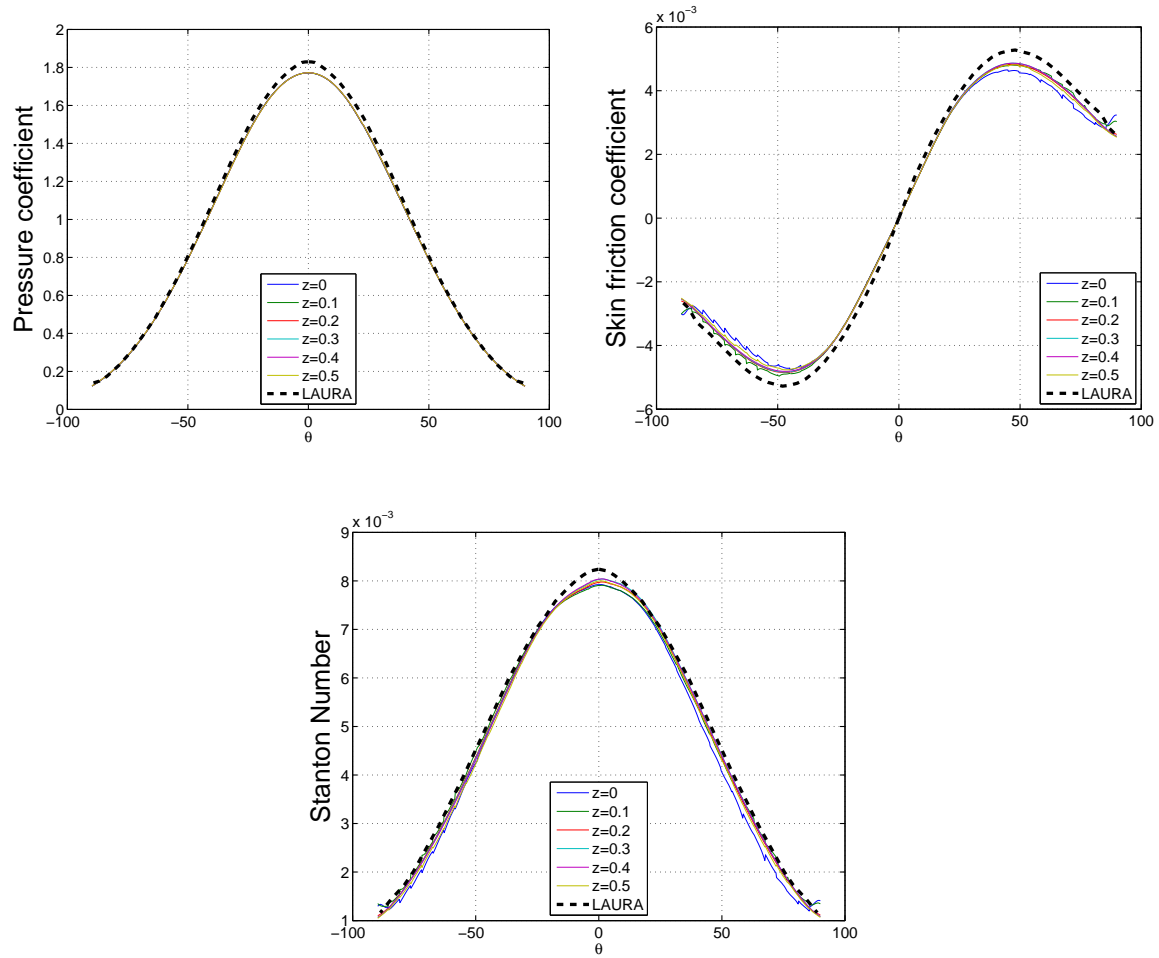


Figure 6-24: Cylinder surface plots of $p = 3$ solution of flow over 3D extruded half cylinder with $M_\infty = 17.605$, $Re = 376,930$.

technique of Oliver [102]. The transformation from x - y coordinates to ξ - η coordinates in the mapped space was,

$$x = \exp(\xi) \cos(\eta); \quad y = \exp(\xi) \sin(\eta) \quad .$$


The transformation therefore unwrapped the curved physical space into a linear computational space. The requested metrics determined by the adaptation framework in physical space were converted to the mapped space. The mesh generation was performed on the transformed metrics by BAMG. The higher-order, cubic geometry nodes were inserted and then all of the node positions were transformed back to physical space. This resulted in cubic geometry elements everywhere in the domain and avoided negative volumes.

The initial and final adapted grids are shown in Figures 6-25a–b and depict the expected refinement of the bow shock, but only to the extent that it impacts the heat flux on the cylinder. The Mach number contours in Figure 6-25c also illustrate the focus of the refinement on the shock in a confined region to create a thin shock layer. The adjoint contours depict the strong sensitivity of the heat load along the cylinder to the stagnation streamline, creating an adjoint wake from the cylinder. This is similarly reflected in the heavy refinement of the stagnation streamline leading up to the boundary layer in the final adapted mesh. Also shown in Figures 6-25e–g is the convergence of the estimated error as a percent of the functional. After eight adaptation iterations, there were approximately three million degrees of freedom and the estimated error was nearly 0.01 percent.


To better visualize the resolution quality of the adapted flow field, Figure 6-26 displays pressure, temperature, and Mach number along the stagnation streamline. For comparison, the same streamline is plotted for the Grid 0 solution of the 2D study. The focus of the adaptation on the shock resolution is quite evident. Additionally, the shock capturing scheme has successfully mitigated all overshoots and undershoots near the hypersonic shock.

The cylinder surface quantities of interest for the adaptation solution are shown in Figure 6-27. In this case, the pressure coefficient matches the LAURA result exactly. This confirms that the mismatch in Figure 6-20 is due to a coarser resolution of the shock. The skin friction displays the same offset from the LAURA results as the previous results. Finally, despite using unstructured meshes everywhere in the domain, for both the shock and boundary layer, there is generally good agreement with LAURA for the predicted heat transfer distribution on the cylinder. However, the heat transfer distribution at the stagnation point exhibits some oscillations. This is due to poor mesh resolution of the boundary layer at the stagnation point in the adapted grid. At the stagnation point, the boundary layer thickness is quite small and requires considerable grid density. The adaptation, not yet fully robust for hypersonic problems, did not sufficiently refine this region of the flowfield for a few reasons. First, the integral output hides the small oscillations in heat transfer from the error estimate. Also, the strong bow shock dominates the error estimate and is the central focus of the grid adaptation. Finally, since the cell anisotropy is determined by derivatives of the Mach number, the elements at the stagnation point have very little stretching.


A modified adaptation sequence was performed to underscore the causes of the noise in




(a) Initial mesh




(b) Final mesh




(c) Mach




(d) Energy adjoint



(e) Adaptation error history



(f) Adaptation functional convergence



(g) Adaptation envelope convergence

Figure 6-25: Initial and final adapted mesh, cylinder contour plots, error history and functional convergence for adaptation of flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$, $p = 2$.




Figure 6-26: Stagnation line plots of $p = 3$ solution of flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$ (taken from Grid 0).

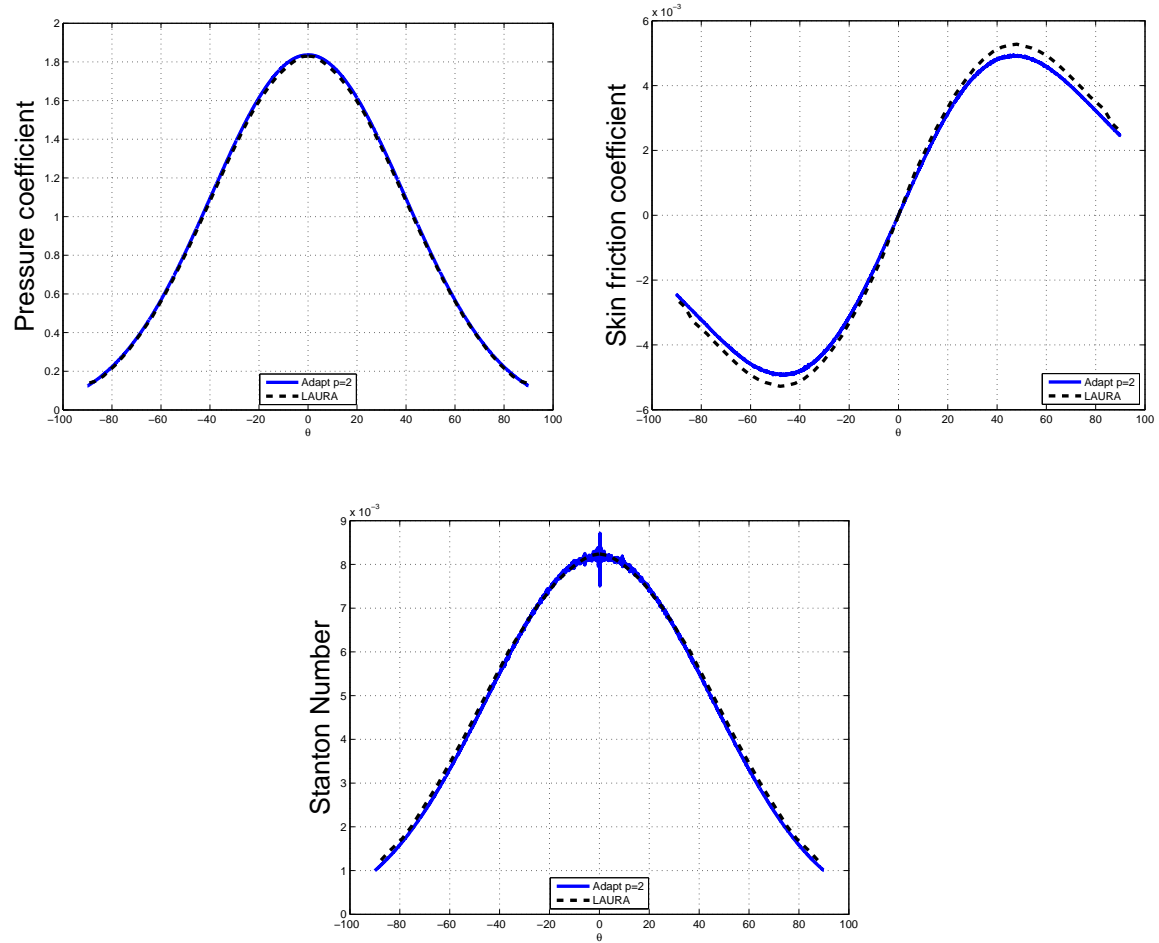


Figure 6-27: Cylinder surface plots of adapted solution of flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$.

the surface heat transfer in Figure 6-27c. This ancillary adaptation sequence was initiated from an intermediate adapted solution and progressed through a few adaptation iterations. The modifications made to the adaptation sequence were twofold. First, to emphasize that the oscillations are hidden from the error estimate, the elemental error contributions were artificially augmented for all of the cells with a centroid located less than 0.01% of the cylinder radius to the wall and within twenty degrees from the stagnation point. Second, to demonstrate that the oscillations are driven by grid errors in the boundary layer and are not convected from the shock, the element metric request used by BAMG to generate a new mesh topology was frozen for all elements with a centroid beyond 1.2 cylinder radii, thereby holding the mesh in the vicinity of the bow shock constant. Using these modifications over a few adaptation iterations generated the surface heat transfer distribution in Figure 6-28a, and successfully eliminated the noise observed in Figure 6-27c. The impact of the modifications on the boundary layer mesh is also depicted in Figure 6-28b, which tracks the average radial distance of the first point off of the cylinder wall within ten degrees of the stagnation point. The modified adaptation sequence clearly refines the boundary layer much more rapidly than the standard adaptation without modification.




Figure 6-28: Cylinder surface plots of adapted solution of flow over 2D half cylinder with $M_\infty = 17.605$, $Re = 376,930$.

Chapter 7

Conclusions

7.1 Summary and Contributions

This thesis has presented a shock capturing methodology for higher-order ($p > 1$) methods. Higher-order schemes are one of the target growth areas in CFD, as they offer advantages and capabilities over traditional second-order accurate methods in computational efficiency. One such use of these methods is the near-field prediction of pressure signatures generated by a supersonic aircraft that can be used in sonic boom propagation models. For this application, adaptive higher-order solutions might enable the design of a supersonic aircraft quiet enough to fly over populated areas. Another application well suited to the use of higher-order methods is the hypersonic flow regime. At hypersonic speeds, nonlinearities and complex physical phenomenon complicate accurate numerical estimation of engineering quantities. Additionally, the significant monetary or human risk involved in reentry flight make high-confidence simulations a valuable asset. While unstructured grids are best suited for the meshing of complex geometries, previous research has shown that surface heating predictions are vulnerable to the variability inherent in unstructured grids. Higher-order solutions might be able to overcome this shortcoming and yield accurate heat transfer estimates. This thesis has demonstrated the following contributions:

- **Motivation for a smooth representation of artificial viscosity for shock capturing in higher-order solutions and a formulation to achieve that representation in the context of the compressible Navier-Stokes equations.**

In higher-order solutions, the strong numerical noise in a shock obstructs the point-wise addition of artificial viscosity. Shock indicators are therefore integral, piecewise-constant functions in the domain. If these shock indicators are used as the basis for artificial viscosity addition, then the artificial viscosity distribution, in its simplest form, would be non-smooth and introduce unwanted errors into the discretization. Specifically, to conserve the viscous flux across element boundaries, a jump in viscosity requires a similar jump in state derivatives. In multiple dimensions, this element-to-element variation, coupled with shock strength changes in a curved bow shock and/or changes in the cell size and orientation, can pollute the downstream flow field. To achieve a higher-order variation of artificial viscosity, this work proposed a new PDE,

to be solved in conjunction with the original governing equations, that governs the distribution of artificial viscosity in a computational domain. This PDE uses existing shock indicators as forcing functions to only apply artificial viscosity in the vicinity of discontinuities. Both inter-element jumps and the decay rate of polynomial expansion coefficients were shown to be reliable shock detection metrics. The PDE was implemented for the compressible Navier-Stokes equations using the discontinuous Galerkin finite element method in a manner that preserves total enthalpy through a shock.

- **Modification of the dual-weighted residual error estimation and adaptation framework for flows with discontinuities and application to supersonic and hypersonic cases.**

The PDE-based artificial viscosity can be used in conjunction with an automated grid adaptation framework based on the error estimation of an output functional. If the error estimate is computed with the adjoint solution using the dual-weighted-residual method of the expanded system of equations, then the additional error contributions from the artificial viscosity are accounted for. This work applied h -adaptation techniques to a series of test cases, culminating in a sample 2D adaptive computation of the drag and far-field pressure signature generated by an airfoil in supersonic flow. All adaptation cases demonstrated a computational efficiency benefit for moving to higher-order solutions ($p > 1$). However, in discontinuous flows, the error in the vicinity of the shock is $\mathcal{O}(h/p)$. Since the degrees of freedom scale with p^d , the marginal improvement in computational efficiency for discontinuous flows decreases for increasing values of p .

- **Demonstration of accurate surface heating, shear stress, and pressure prediction for hypersonic problems using unstructured and adapted grids.**

The PDE-based artificial viscosity was also successfully applied to hypersonic flow cases. The first test case was the flow over a 15° compression ramp. Solutions on a series of structured meshes designed to achieve grid convergence agreed well with those of other computational codes, but not experimental data. It was also shown that adaptation could be used to achieve the same results as a well designed structured mesh, despite robustness challenges for hypersonic adaptation. Another test case that was examined in depth was the hypersonic flow over a half-cylinder in both 2D and 3D. Using a series of hybrid structured-unstructured meshes (structured meshes for the boundary layer and unstructured meshes for the external flow), the higher-order solutions with the PDE-based artificial viscosity demonstrated good prediction of surface heating and were less susceptible to the errors introduced by the unstructured grid. Furthermore, allowing the adaptation framework to modify the grid for minimal error in the integrated heat load to the cylinder also resulted in the same heat transfer distribution as a structured mesh.

7.2 Future Work

- **Further application testing of the PDE-based artificial viscosity.**

In this work, the PDE-based artificial viscosity model has demonstrated advantages over a non-smooth representation for higher-order DG solutions. However, the test cases contained in this thesis were limited to a few applications. Within the context of the compressible Navier-Stokes equations, PDE-based artificial viscosity should be explored for unsteady flow problems, turbulent flows requiring the use of the RANS, and a greater selection of more complex transonic, supersonic and hypersonic flow cases. Moreover, the PDE-based artificial viscosity can be applied to other equation sets, such as the magnetohydrodynamics (MHD) equations.

- **Extension of the PDE-based artificial viscosity to other discretizations.**

In addition to the further application of the PDE-based artificial viscosity to other test cases and equation sets, it should also be applied to other discretizations as well. The work in this thesis has focused on the discontinuous Galerkin finite element method. However, the benefits of a smooth variation in artificial viscosity should apply to other higher-order discretizations as well. Additional study on the behavior and benefits of a smooth variation of viscosity in a finite volume and continuous finite element setting is one avenue of future research.

- **Extension and improvement of the adaptation mechanics**

The adaptation cases presented in this work were strictly two-dimensional. No 3D adaptation cases were pursued due to the need for unstructured, anisotropic, metric-based 3D meshing. Additionally, the adaptation framework focused solely on h -adaptation. Since the errors in the vicinity of the shock scale with $\mathcal{O}(h/p)$ while the degrees of freedom scale with $\mathcal{O}(p^d)$, the computational efficiency benefit of higher-order interpolations drops off for higher and higher values of p . Adding an hp -adaptive capability should allow for better degree of freedom management and optimization in the adaptation process. In smooth flow regions, refinement could be done with p and near discontinuities refinement could be done with h . In addition, the 2D test case of a supersonic airfoil demonstrated that adaptation for drag or the far-field pressure signature place very different demands on the mesh. A multi-objective grid adaptation strategy would be better suited for a multi-disciplinary design setting employing CFD analysis. Finally, although the adapted solutions for the hypersonic applications arrived at the same results as well-designed structured meshes, there are shortcomings in the robustness of the adaptation mechanics for the hypersonic flow regime.

- **Use in gradient-based design optimization framework.**

One of the motivating factors for the use of artificial viscosity is that it is well suited to adjoint-based computational tools, such as calculation of design variable sensitivities in gradient driven optimization and in output-based error estimation. While variable sensitivities calculated via the adjoint were mentioned briefly in Chapter 3, much more attention was given to output-based error estimation and its link to automated grid adaptation. Further work is needed to explore the behavior of the PDE-based

artificial viscosity in the context of optimization methods that use the adjoint to compute sensitivity derivatives.

- **Other avenues towards smooth artificial viscosity.**

Chapter 3 presented motivating factors from the one-dimensional Burgers equation and the multi-dimensional Navier-Stokes equations for a smooth variation of viscosity. This work proposed an artificial viscosity that was governed by an elliptic PDE as a means to achieve a smooth variation, as it maintained the compact DG stencil and relied upon the existing shock detection algorithms. One obvious drawback to this method is the augmented state vector and increased computational load. There might be other approaches by which one could arrive at a smooth variation of viscosity for higher-order solutions, while using fewer degrees of freedom and still maintain a compact stencil.

Appendix A

Dual Consistency of Nonlinear Viscosity

Just as consistency implies that the exact primal solution satisfies the discrete primal residual, dual consistency implies that the exact dual solution satisfies the discrete dual residual. However, consistency of the primal problem does not necessarily ensure consistency in the dual problem [103]. A dual consistent discretization is important for adjoint-based analysis methods. For instance, accurate output-based error estimation relies upon an accurate adjoint solution, which can be corrupted for dual inconsistent discretizations [63, 87]. This chapter briefly touches on the implications of shock capturing on the dual consistency of the discretization presented in Chapter 2.

A.1 Dual Consistency Preliminaries

Let $u \in \mathcal{V}$, where \mathcal{V} is a given function space and $\mathcal{L} : \mathcal{V} \rightarrow \mathbb{R}$ be the linear differential operator of the equation $\mathcal{L}u = f$ in a domain, $\Omega \subset \mathbb{R}^d$, and $f \in L^2(\Omega)$. The adjoint, $\psi \in \mathcal{V}$, for a given linear functional, $\mathcal{J} : \mathcal{V} \rightarrow \mathbb{R}$, is determined by the linear dual problem,

$$\mathcal{L}^*\psi = \mathcal{J}, \quad \text{where} \quad \langle \mathcal{L}v, w \rangle = \langle v, \mathcal{L}^*w \rangle, \quad (\text{A.1})$$

and $\langle \cdot, \cdot \rangle$ denotes the inner product. In the interest of simplicity, the role of boundary conditions upon the adjoint solution and dual consistency is ignored. A more thorough presentation on the role of boundary conditions can be found in [109] and [87].

If \mathcal{L} is instead a non-linear differential operator for the equation $\mathcal{L}(u) = f$, then the adjoint is determined by the linearized dual problem,

$$\mathcal{L}'[u](\psi) = \mathcal{J}'[u], \quad \text{where} \quad \langle \mathcal{L}'[u](v), w \rangle = \langle v, \mathcal{L}'[u](w) \rangle, \quad (\text{A.2})$$

and $\mathcal{L}'[u]$ denotes the Frechet derivative at u ,

$$\mathcal{L}'[u](v) \equiv \lim_{\sigma \rightarrow 0} \frac{\mathcal{L}(u + \sigma v) - \mathcal{L}(u)}{\sigma}.$$

Consistency and dual consistency can be determined by applying the solution, u , and the dual solution, ψ , to the numerical discretization. Define the typical DG space, \mathcal{V}_H^p , to be the finite vector space of piecewise-polynomial functions of degree p on every element, κ , within the triangulation, \mathcal{T}_h , of the domain, $\Omega = \bigcup_{\kappa \in \mathcal{T}_h} \kappa$,

$$\mathcal{V}_H^p \equiv \{v \in L^2(\Omega) \mid v|_{\kappa} \in P^p, \forall \kappa \in \mathcal{T}_h\},$$

Furthermore, let $\mathcal{W}_H^p \equiv \mathcal{V}_H^p + \mathcal{V}$, the sum of the continuous and discrete function spaces.

Let the discrete primal and dual solutions, $u_H \in \mathcal{V}_H^p$ and $\psi_H \in \mathcal{V}_H^p$, satisfy,

$$\begin{aligned} \mathcal{R}(u_H, v_H) &= 0, \quad \forall v_H \in \mathcal{V}_H^p, \\ \mathcal{R}'[u_H](v_H, \psi_H) &= \mathcal{J}'_H[u_H](v_H), \quad \forall v_H \in \mathcal{V}_H^p, \end{aligned}$$

where $\mathcal{R} : \mathcal{W}_H^p \times \mathcal{W}_H^p \rightarrow \mathbb{R}$ is a semi-linear form (linear in the second argument) and the weak discretization of $\mathcal{L}(u)$. \mathcal{R} is said to be consistent if, given the exact solution $u \in \mathcal{V}$,

$$\mathcal{R}(u, v_H) = \mathcal{R}(u_H, v_H), \quad \forall v_H \in \mathcal{V}_H^p.$$

Meaning that the exact, continuous solution satisfies the discrete residual. Similarly, the discretization is declared dual consistent if, given the exact solution $\psi \in \mathcal{V}$,

$$\mathcal{R}'[u](v, \psi) = \mathcal{J}'_H[u](v), \quad \forall v \in \mathcal{W}_H^p. \quad (\text{A.3})$$

A.2 Dual Consistency of the Non-Linear Poisson Equation

The use of shock capturing with artificial viscosity creates a non-linear dependence of the diffusion on the conservative state vector. The implications of this dependence, specifically for the discretization in Chapter 2, are explored by examining the Poisson equation with a non-linear viscosity and a non-linear source term,

$$\begin{aligned} -\nabla \cdot (\nu \nabla u) - f &= 0 \quad \text{in } \Omega \subset \mathbb{R}^d, \\ u &= 0 \quad \text{on } \partial\Omega, \end{aligned} \quad (\text{A.4})$$

where $u \in \mathcal{V} \equiv H^2(\Omega)$, $\nu \in \mathcal{C}^1(\mathbb{R}^{d+1})$ and $f \in \mathcal{C}^1(\mathbb{R}^{d+1})$. Here $\nu = \nu(u, \nabla u)$ is the non-linear viscosity and represents the shock indicator in the non-smooth formulation of artificial viscosity, (4.1). The source term is $f = f(u, \nabla u)$ and represents the non-linear shock indicator appearing in the artificial viscosity equation (4.2).

Due to the non-linearity of (A.4), it is not self-adjoint, and the adjoint solution, $\psi \in H^2(\Omega)$, is given by the dual problem through Frechet differentiation,

$$\begin{aligned} -\nabla \cdot (\nu \nabla \psi) + \nabla \psi \cdot \nu_u \nabla u - \nabla \cdot (\nabla \psi \cdot \nu_{\nabla u} \nabla u) - \\ f_u \psi + \nabla \cdot (\psi f_{\nabla u}) - \mathcal{J}'[u] &= 0 \quad \text{in } \Omega, \end{aligned}$$

where the notation ν_u denotes the variation of ν with respect to u .

The standard variational formulation of (A.4) involves the multiplication by a test function and integration by parts. Before this is done, however, (A.4) is written as a first-order system,

$$\begin{aligned} -\nabla \cdot \mathcal{Q} - f &= 0 & \text{in } \Omega, \\ \mathcal{Q} - \nu \nabla u &= 0 & \text{in } \Omega, \\ u &= 0 & \text{on } \partial\Omega, \end{aligned}$$

and the DG discretization is derived as in Section 2.1. For the purposes of this analysis, the notation of Arnold et al. [5] is adopted, who present all prominent elliptic discretizations for DG, including BR2, in a unified form:

$$\mathcal{R}(u_H, v_H) \equiv B_H(u_H, v_H) - \int_{\Omega} v_H f \, dx, \quad (\text{A.5})$$

where

$$\begin{aligned} B_H(u_H, v_H) \equiv & \int_{\Omega} \nabla v_H \cdot \nu \nabla u_H \, dx + \int_{\Gamma} \left(\llbracket \hat{u} - u_H \rrbracket \cdot \{\nu \nabla v_H\} - \llbracket v_H \rrbracket \cdot \{\hat{\mathcal{Q}}\} \right) ds \\ & + \int_{\Gamma^0} \left(\{\hat{u} - u_H\} \llbracket \nu \nabla v_H \rrbracket - \{v_H\} \llbracket \hat{\mathcal{Q}} \rrbracket \right) ds, \quad (\text{A.6}) \end{aligned}$$

\hat{u} and $\hat{\mathcal{Q}}$ are the numerical approximations to the state and viscous flux along the discontinuous element edges, Γ is the union of element boundaries in the triangulation and $\Gamma^0 = \Gamma / \partial\Omega$ (interior faces only).

Arnold et al. investigated the dual consistency of (A.5) for the linear Poisson equation, $\nabla^2 u = f$. They determined that if the numerical fluxes, \hat{u} and $\hat{\mathcal{Q}}$, are both *conservative* (meaning that the fluxes are single-valued along element edges) and *consistent* (implying that $\hat{u}|_{\Gamma} = u|_{\Gamma}$), then the discretization in (A.5) is dual consistent. Additionally, Oliver and Darmofal [103] specifically addressed the issue of a non-linear source term dependent on the state and state gradients and discretized in DG by (A.5). They prove that if the source term is a function of the state gradients, $f = f(u, \nabla u)$, then (A.5) is dual inconsistent. The analysis below of the non-linear Poisson equation follows the methods used by Oliver and Darmofal.

To ascertain the dual consistency of the discretization, one must first determine, $\mathcal{R}'[u](v, \psi)$, which is done by separating out contributions from the diffusion and source terms,

$$\begin{aligned} \mathcal{R}'[u](v_H, \psi) &= B'_H[u](v_H, \psi) - \int_{\Omega} v_H f'[u] \, dx \\ &= B'_H[u](v_H, \psi) - \int_{\Omega} v_H [f_u \psi - \nabla \cdot (f_{\nabla u} \psi)] \, dx \\ &\quad - \int_{\Gamma} \llbracket v_H \rrbracket \cdot \{f_{\nabla u} \psi\} \, ds - \int_{\Gamma^0} \{v_H\} \llbracket f_{\nabla u} \psi \rrbracket \, ds. \end{aligned}$$

The linearization of B_H is,

$$\begin{aligned}
B'_H[u](v_H, \psi) &= \int_{\Omega} \nabla \psi \cdot \left[\nu_u v_H \nabla u + (\nu_{\nabla u} \cdot \nabla v_H) \nabla u + \nu \nabla v_H \right] dx + \\
&\int_{\Gamma} [\![\hat{u}(v) - v_H]\!] \cdot \{\nu \nabla \psi\} + [\![\hat{u} - u]\!] \cdot \{(\nu_u v_H + \nu_{\nabla u} \cdot \nabla v_H) \nabla \psi\} - [\![\psi]\!] \cdot \{\hat{\mathcal{Q}}'[u]\} ds + \\
&\int_{\Gamma^0} \{\hat{u}(v) - v_H\} \cdot [\![\nu \nabla \psi]\!] + \{\hat{u} - u\} \cdot [\![\nu_u v_H + \nu_{\nabla u} \cdot \nabla v_H] \nabla \psi]\!] - \{\psi\} \cdot [\![\hat{\mathcal{Q}}'[u]\!] ds. \quad (\text{A.7})
\end{aligned}$$

Integrating by parts, one obtains the following identities,

$$\begin{aligned}
\int_{\Omega} \nabla \psi \cdot \nu \nabla v_H &= - \int_{\Omega} v_H \nabla \cdot (\nu \nabla \psi) dx + \int_{\Gamma} [\![v_H]\!] \cdot \{\nu \nabla \psi\} ds \\
&\quad + \int_{\Gamma^0} \{v_H\} [\![\nu \nabla \psi]\!] ds, \\
\int_{\Omega} \nabla \psi \cdot (\nu_{\nabla u} \cdot \nabla v_H) \nabla u &= - \int_{\Omega} v_H \nabla \cdot (\nabla \psi \cdot \nu_{\nabla u} \nabla u) dx \\
&\quad + \int_{\Gamma} [\![v_H]\!] \cdot \{\nabla \psi \cdot \nu_{\nabla u} \nabla u\} ds \\
&\quad + \int_{\Gamma^0} \{v_H\} [\![\nabla \psi \cdot \nu_{\nabla u} \nabla u]\!] ds.
\end{aligned}$$

These expressions are substituted back into (A.7). Additionally, if u and ψ solve the continuous primal and dual problems, then $u \in H^2(\mathcal{T}_H)$ and $\psi \in H^2(\mathcal{T}_H)$. Therefore, $\{\psi\} = \psi$, $\{\nabla \psi\} = \nabla \psi$, $[\![\psi]\!] = 0$ and $[\![\nabla \psi]\!] = 0$ (with similar assumptions for u). Assuming a conservative scheme, then the jump in fluxes at element boundaries is zero, $[\![\hat{u}]\!] = [\![\hat{\mathcal{Q}}]\!] = 0$. Furthermore, a consistent scheme implies that $\{\hat{u}\}|_{\Gamma} - u|_{\Gamma} = 0$. These simplifications yield,

$$\mathcal{R}'[u](v_H, \psi) - \mathcal{J}'_H[u](v_H) = \int_{\Gamma} [\![v_H]\!] \cdot [(\nabla \psi \cdot \nu_{\nabla u} \nabla u) - (f_{\nabla u} \psi)] ds \quad . \quad (\text{A.8})$$

Thus, the unified discretization for elliptic operators, as defined by Arnold et al., is *dual inconsistent* if the viscosity or source term are functions of both the state and gradient of the state. This supports the findings of Oliver and Darmofal [103]. One can see from the analysis that it is specifically the gradient of the state that leads to the inconsistency.

To achieve dual consistency for the non-linear Poisson problem, one would need to use a different discretization than the one presented by Arnold et al., which includes the BR2 discretization described in Section 2.1.

The dual consistency analysis also has implications for the selection of a shock indicator. If the shock indicator is a function of state gradients, then the discretization of Section 2.1 for both the non-smooth and PDE-based artificial viscosity formulations will be dual inconsistent. The resolution indicator involves the restriction of a function from order p to order $p - 1$, a derivative-like operation, and is therefore dependent on state gradients. The jump indicator, however, is a function of the state only, but can only be used with the PDE-based artificial viscosity to keep the numerical stencil compact.

Bibliography

- [1] Adjerid, S., Devine, K. D., Flaherty, J. E., and Krivodonova, L. “A posteriori error estimation for discontinuous Galerkin solutions of hyperbolic problems.” *Comput. Methods Appl. Mech. Eng.*, 191:1097–1112, 2002.
- [2] Aliabadi, S., Tu, S.-Z., and Watts, M. “An alternative to limiter in discontinuous Galerkin finite element method for simulation of compressible flows.” AIAA Paper 2004-76, 2004.
- [3] Allmaras, S. R. *A coupled Euler/Navier-Stokes algorithm for 2-D unsteady transonic shock/boundary-layer interaction*. PhD thesis, Massachusetts Institute of Technology, 1989.
- [4] Allmaras, S. R. and Giles, M. B. “A second-order flux split scheme for the unsteady 2-D Euler equations on arbitrary meshes.” AIAA Paper 87-1119-CP, 1987.
- [5] Arnold, D., Brezzi, F., Cockburn, B., and Marini, L. “Unified Analysis of Discontinuous Galerkin Methods for Elliptic Problems.” *SIAM J. Numer. Anal.*, 39(5):1749–1779, 2002.
- [6] Atkins, H. L. “Continued Development of the Discontinuous Galerkin Method for Computational Aeroacoustic Applications.” AIAA Paper 97-1581, 1997.
- [7] Babuska, I. and Miller, A. “The Post-Processing Approach in the Finite Element Method- Part 1: Calculation of Displacements, Stresses and Other Higher Derivatives of the Displacements.” *International Journal for Numerical Methods in Engineering*, 20:1085–1109, 1984.
- [8] Babuska, I. and Miller, A. “The Post-Processing Approach in the Finite Element Method- Part 2: The Calculation of Stress Intensity Factors.” *International Journal for Numerical Methods in Engineering*, 20:1111–1129, 1984.
- [9] Babuska, I. and Miller, A. “The Post-Processing Approach in the Finite Element Method- Part 3: A Posteriori Error Estimates and Adaptive Mesh Selection.” *International Journal for Numerical Methods in Engineering*, 20:2311–2324, 1984.
- [10] Baldwin, B. S. and MacCormack, R. W. “Interaction of Strong-Shock Wave with Turbulent Boundary Layer.” AIAA Paper 74-558, 1974.
- [11] Barth, T. J. “Numerical methods for gasdynamic systems on unstructured meshes.” In Rohde, C., Kroner, D., and Ohlberger, M., editors, *An Introduction to Recent Developments in Theory and Numerics for Conservation Laws, Lecture Notes in Computational Science and Engineering, Vol. 5*, pages 195–285. Springer, Berlin, 1997.
- [12] Barth, T. and Larson, M. “A posteriori error estimates for higher order Godunov finite volume methods on unstructured meshes.” In Herban, R. and Kröner, D., editors, *Finite Volumes for Complex Applications III*, London, 2002. Hermes Penton.
- [13] Bassi, F. and Rebay, S. “Accurate 2D Euler Computations by means of a High Order Discontinuous Finite Element Method.” In Deshpande, S. M., Desai, S. S., and Narasimha, R., editors, *Fourteenth International Conference on Numerical Methods in Fluid Dynamics, Lecture Notes in Physics vol. 453*, pages 234–239. Springer, Berlin, 1994.

- [14] Bassi, F. and Rebay, S. “A High-Order Accurate Discontinuous Finite Element Method for the Numerical Solution of the Compressible Navier-Stokes Equations.” *J. Comput. Phys.*, 131: 267–279, 1997.
- [15] Bassi, F. and Rebay, S. “High-Order Accurate Discontinuous Finite Element Solution of the 2-D Euler Equations.” *J. Comput. Phys.*, 138:251–285, 1997.
- [16] Bassi, F. and Rebay, S. “GMRES Discontinuous Galerkin Solution of the Compressible Navier-Stokes Equations.” In Cockburn, K. and Shu, editors, *Discontinuous Galerkin Methods: Theory, Computation and Applications*, pages 197–208. Springer, Berlin, 2000.
- [17] Becker, R. and Rannacher, R. “An optimal control approach to a posteriori error estimation in finite element methods.” In Iserles, A., editor, *Acta Numerica*. Cambridge University Press, 2001.
- [18] Bey, K. and Oden, J. T. “A Runge-Kutta discontinuous finite element method for high speed flows.” AIAA Paper 91-1575, June 1991.
- [19] Bey, K. and Oden, J. T. “hp-version discontinuous Galerkin methods for hyperbolic conservation laws.” *Comput. Methods Appl. Mech. Engrg.*, 133:259–286, 1996.
- [20] Bilardo, V. J., Curran, F. M., Hunt, J. L., Lovell, N. T., Maggio, G., Wilhite, A. W., and McKinney, L. E. “The Benefits of Hypersonic Airbreathing Launch Systems for Access to Space.” AIAA Paper 2003-5265, 2003.
- [21] Brachet, J.-B. “A Dynamic Multiscale Viscosity Algorithm for Shock Capturing in Runge Kutta Discontinuous Galerkin Methods.” Master’s thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2005.
- [22] Burbeau, A., Sagaut, P., and Bruneau, C.-H. “A problem-independent limiter for high-order Runge-Kutta discontinuous Galerkin methods.” *Journal of Computational Physics*, 169(1): 111–150, 2001.
- [23] Candler, G. V. and Nompelis, I. “CFD Validation for Hypersonic Flight: Real Gas Flows.” AIAA Paper 2002-0434, 2002.
- [24] Castro-Díaz, M. J., Hecht, F., Mohammadi, B., and Pironneau, O. “Anisotropic Unstructured Mesh Adaption for Flow Simulations.” *Int. J. Numer. Meth. Fluids*, 25:475–491, 1997.
- [25] Chan, M. K. *Supersonic Aircraft Optimization for Minimizing Drag and Sonic Boom*. PhD thesis, Stanford University, Stanford, CA, 2003.
- [26] Chen, G.-Q., Du, Q., and Tadmor, E. “Spectral Viscosity Approximations to Multidimensional Scalar Conservation Laws.” *Mathematics of Computation*, 61:629–643, 1993.
- [27] Choi, S., Alonso, J. J., and Van der Weide, E. “Numerical and Mesh Resolution Requirements for Accurate Sonic Boom Prediction of Complete Aircraft Configurations.” AIAA Paper 2004-1060, 2004.
- [28] Cockburn, B., Hou, S., , and Shu, C. W. “TVB Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws IV: The multidimensional case.” *Math. Comp.*, 54:545–581, 1990.
- [29] Cockburn, B., Karniadakis, G. E., and Shu, C. W. “The Development of discontinuous Galerkin methods.” In Cockburn, K. and Shu, editors, *Discontinuous Galerkin Methods: Theory, Computation and Applications*, pages 3–50. Springer, Berlin, 2000.

- [30] Cockburn, B., Lin, S. Y., and Shu, C. W. “TVB Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws III: One dimensional systems.” *J. Comput. Phys.*, 84:90–113, 1989.
- [31] Cockburn, B. and Shu, C. W. “TVB Runge-Kutta local projection discontinuous Galerkin finite element method for scalar conservation laws II: General framework.” *Math. Comp.*, 52: 411–435, 1989.
- [32] Cockburn, B. and Shu, C. W. “The Runge-Kutta local projection P^1 -discontinuous Galerkin method for scalar conservation laws.” *RAIRO Model Math. Anal. Numer.*, 25:337–361, 1991.
- [33] Cockburn, B. and Shu, C. W. “The Local Discontinuous Galerkin Method for Time-Dependent Convection-Diffusion Systems.” *SIAM J. Numer. Anal.*, 35(6):2440–2463, December 1998.
- [34] Cockburn, B. and Shu, C. W. “The Runge-Kutta discontinuous Galerkin finite element method for conservation laws V: Multidimensional systems.” *J. Comput. Phys.*, 141:199–224, 1998.
- [35] Cockburn, B. and Shu, C. W. “Runge-Kutta Discontinuous Galerkin Methods for Convection-Dominated Problems.” *J. of Sci. Comput.*, 16(3):173–261, September 2001.
- [36] Coen, P., Long-Davis, M. J., and Povinelli, L. “Fundamental Aeronautics Program Supersonics Project Reference Document.” NASA, 2006. <http://www.aeronautics.nasa.gov/fap/documents.html>.
- [37] Committee on Breakthrough Technology for Commercial Supersonic Aircraft. *Commercial Supersonic Technology, The Way Ahead*. National Academies Press, Washington, D.C., 2001.
- [38] Dafermos, C. M. *Hyperbolic Conservation Laws in Continuum Physics*. Springer, Berlin, 2005.
- [39] Diosady, L. T. “A Linear Multigrid Preconditioner for the Solution of the Navier-Stokes Equations using a Discontinuous Galerkin Discretization.” Masters thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, May 2007.
- [40] Dolejsi, V., Feistauer, M., and Schwab, C. “On some aspects of the discontinuous Galerkin finite element method for conservation laws.” *Mathematics and Computers in Simulation*, 61 (3):333–346, 2003.
- [41] Fedkiw, R. P., Sapiro, G., and Shu, C.-W. “Shock capturing, level sets, and PDE based methods in computer vision and image processing - A review of Osher’s contributions.” *Journal of Computational Physics*, 185(2):309–341, 2003.
- [42] Fidkowski, K. “A High-Order Discontinuous Galerkin multigrid solver for aerodynamic applications.” Master’s thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2004.
- [43] Fidkowski, K. J. *A Simplex Cut-Cell Adaptive Method for High-Order Discretizations of the Compressible Navier-Stokes Equations*. PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2007.
- [44] Fidkowski, K. J., Oliver, T. A., Lu, J., and Darmofal, D. L. “ p -Multigrid Solution of High-Order Discontinuous Galerkin Discretizations of the Compressible Navier-Stokes Equations.” *Journal of Computational Physics*, To appear:1–2, 2005.
- [45] Gelb, A. and Tadmor, E. “Enhanced Spectral Viscosity Approximations for Conservation Laws.” *Applied Numerical Mathematics*, 33:3–21, 2000.
- [46] Giles, M. B. and Süli, E. “Adjoint Methods for PDEs: *a posteriori* Error Analysis and Postprocessing by Duality.” *Acta Numerica*. Cambridge University Press, 2002.

- [47] Giles, M. “Discrete Adjoint Approximations with Shocks.” In Hou, T. and Tadmor, E., editors, *Hyperbolic Problems: Theory, Numerics, Applications*. Springer-Verlag, 2003.
- [48] Giles, M. B. “Analysis of the Accuracy of Shock-Capturing in the Steady Quasi-1D Euler Equations.” *Computational Fluid Dynamics Journal*, 5:247–258, 1996.
- [49] Giles, M. B. and Pierce, N. A. “An Introduction to the Adjoint Approach to Design.” *Flow, Turbulence and Combustion*, 65:393–415, 2000.
- [50] Giles, M. B., Pierce, N. A., and Suli, E. “Progress in adjoint error correction for integral functionals.” *Computing and Visualization in Science*, 6:113–121, 2004.
- [51] Gnoffo, P. A. “Computational Fluid Dynamics Technology for Hypersonic Applications.” AIAA Paper 2003-2829, 2003.
- [52] Gnoffo, P. A. and Cheatwood, F. M. *User’s Manual for the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA)*. NASA Langley Research Center, Virginia, 1996. NASA Technical Memorandum 4674.
- [53] Gnoffo, P. A., Gupta, R. N., and Shinn, J. L. “Conservation Equations and Physical Models for Hypersonic Air Flows in Thermal and Chemical Nonequilibrium.” NASA Technical Paper 2867, 1989.
- [54] Gnoffo, P. A., Weilmuenster, K. J., Hamilton, H. H., Olynick, D. R., and Venkatapathy, E. “Computational Aerothermodynamic Design Issues for Hypersonic Vehicles.” *Journal of Spacecraft and Rockets*, 36:21–43, 1999.
- [55] Gnoffo, P. A. and White, J. A. “Computational Aerothermodynamic Simulation Issues for Unstructured Grids.” AIAA Paper 2004-2371, 2004.
- [56] Guermond, J. L. “A Finite Element Technique for Solving First-Order PDEs in L^{p*} .” *SIAM Journal of Numerical Analysis*, 42(2):714–737, 2004.
- [57] Guo, B.-Y., Ma, H.-P., and Tadmor, E. “Spectral Vanishing Viscosity Method for Nonlinear Conservation Laws.” *SIAM Journal on Numerical Analysis*, 39(4):1254–1268, 2001.
- [58] Habashi, W. G., Dompierre, J., Bourgault, Y., Ait-Ali-Yahia, D., Fortin, M., and Vallet, M.-G. “Anisotropic mesh adaptation: towards user-independent, mesh-independent and solver-independent CFD. Part I: general principles.” *Int. J. Numer. Meth. Fluids*, 32:725–744, 2000.
- [59] Hänel, D., Schwane, R., and Seider, G. “On the Accuracy of Upwind Schemes for the Solution of the Navier-Stokes Equations.” AIAA Paper 87-1105, 1987.
- [60] Harten, A., Engquist, B., Osher, S., and Chakravarthy, S. “Uniformly high order accurate essentially non-oscillatory schemes. III.” *Journal of Computational Physics*, 71:231–303, 1987.
- [61] Harten, A. and Osher, S. “Uniformly high order accurate non-oscillatory schemes. I.” *SIAM Journal on Numerical Analysis*, 24(2):279–309, 1987.
- [62] Hartmann, R. “Adaptive discontinuous Galerkin methods with shock-capturing for the compressible Navier-Stokes Equations.” *Int. J. Numer. Meth. Fluids*, 51:1131–1156, 2006.
- [63] Hartmann, R. and Houston, P. “Adaptive Discontinuous Galerkin Finite Element Methods for the Compressible Euler Equations.” *J. Comput. Phys.*, 183:508–532, 2002.
- [64] Hecht, F. *BAMG: Bidimensional Anisotropic Mesh Generator*. INRIA, Rocquencourt, France, 1998. <http://www-rocq1.inria.fr/gamma/cdrom/www/bamg/eng.htm>.

- [65] Holden, M. S. “A Study of Flow Separation in Regions of Shock Wave-Boundary Layer Interaction in Hypersonic Flow.” AIAA Paper 78-1169, 1978.
- [66] Houston, P. and Süli, E. “*hp*-Adaptive Discontinuous Galerkin Finite Element Methods for First-Order Hyperbolic Problems.” *SIAM J. Sci. Comput.*, 23(4):1226–1252, 2001.
- [67] Houston, P. and Süli, E. “A note on the design of *hp*-adaptive finite element methods for elliptic partial differential equations.” *Comput. Methods Appl. Mech. Engrg.*, 194:229–243, 2005.
- [68] Hughes, T. J. R., Franca, L., and Mallet, M. “A new finite element formulation for computational fluid dynamics: I Symmetric forms of the compressible Euler and Navier-Stokes equations and the second law of thermodynamics.” *Comput. Methods Appl. Mech. Engrg.*, 54: 223–234, 1986.
- [69] Hughes, T. J. R. and Mallet, M. “A new finite element formulation for computational fluid dynamics: III The generalized streamline operator for multidimensional advective-diffusive systems.” *Comput. Methods Appl. Mech. Engrg.*, 58:305–328, 1986.
- [70] Hughes, T. J. R. and Mallet, M. “A new finite element formulation for computational fluid dynamics: IV A discontinuity capturing operator for multidimensional advective-diffusive systems.” *Comput. Methods Appl. Mech. Engrg.*, 58:329–336, 1986.
- [71] Hughes, T. J. R., Mallet, M., and Mizukami, A. “A new finite element formulation for computational fluid dynamics: II Beyond SUPG.” *Comput. Methods Appl. Mech. Engrg.*, 54: 341–355, 1986.
- [72] Jameson, A. “Solution of the Euler equations for two-dimensional transonic flow by a multigrid method.” *Applied Mathematics and Computation*, 13:327–356, 1983.
- [73] Jameson, A. “Aerodynamic design via control theory.” *Journal of Science and Computation*, 3:233–260, 1988.
- [74] Jameson, A., Schmidt, W., and Turkel, E. “Numerical simulation of the Euler equations by finite volume methods using Runge-Kutta time stepping schemes.” AIAA Paper 81-1259, 1981.
- [75] Jameson, A. “Analysis and Design of Numerical Schemes for Gas Dynamics, 2: Artificial Diffusion and Discrete Shock Structure.” *International Journal of Computational Fluid Dynamics*, 5:1–38, 1995.
- [76] John D. Anderson, J. *Hypersonic and High Temperature Gas Dynamics*. McGraw-Hill, New York, 1989.
- [77] John D. Anderson, J. *Modern Compressible Flow*. McGraw-Hill, Boston, MA, 1990.
- [78] John D. Anderson, J. *Computational Fluid Dynamics*. McGraw-Hill, Inc., New York, NY, 1995.
- [79] Johnson, C., Szepessy, A., and Hansbo, P. “On the convergence of shock-capturing streamline diffusion finite element methods for hyperbolic conservation laws.” *Math. Comp.*, 54(189): 107–129, 1990.
- [80] Kandil, O. A., Yang, Z., and Bobbitt, P. J. “Prediction of Sonic Boom Signature Using Euler-Full Potential CFD with Grid Adaptation and Shock Fitting.” AIAA Paper 2002-2542, 2002.
- [81] Kirk, B. S. *Adaptive Finite Element Simulation of Flow and Transport Applications on Parallel Computers*. PhD thesis, University of Texas, 2007.

- [82] Krivodonova, L., Xin, J., Remacle, J.-F., Chevaugeon, N., and Flaherty, J. E. “Shock detection and limiting with discontinuous Galerkin methods for hyperbolic conservation laws.” *Applied Numerical Mathematics*, 48(3):323–338, 2004.
- [83] LeVeque, R. J. *Numerical Methods for Conservation Laws*. Birkhäuser Verlag, Basel, 1992.
- [84] Lillard, R. P. and Dries, K. M. “Laminar Heating Validation of the OVERFLOW Code.” AIAA Paper 2005-689, 2005.
- [85] Lindquist, D. R. and Giles, M. B. “Validity of Linearized Unsteady Euler Equations with Shock Capturing.” *AIAA Journal*, 32:46–53, 1994.
- [86] Liu, X.-D., Osher, S., and Chan, T. “Weighted essentially non-oscillatory schemes.” *Journal of Computational Physics*, 115:200–212, 1994.
- [87] Lu, J. *An a Posteriori Error Control Framework for Adaptive Precision Optimization Using Discontinuous Galerkin Finite Element Method*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2005.
- [88] Luo, H., Baum, J. D., and Löhner, R. “A p-multigrid discontinuous Galerkin method for the Euler equations on unstructured grids: application to compressible flows.” *J. Comput. Phys.*, 211:767–783, 2006.
- [89] Luo, H., Baum, J. D., and Löhner, R. “A Hermite WENO-based limiter for discontinuous Galerkin method on unstructured grids.” *J. Comput. Phys.*, 225:686–713, 2007.
- [90] Lytton, C. C. “Solution of the Euler Equations for Transonic Flow over a Lifting Airfoil- The Bernoulli Formulation (Roe/Lytton Method).” *J. Comput. Phys.*, 73:395–431, 1987.
- [91] Machiels, L., Peraire, J., and Patera, A. T. “A *a posteriori* Finite Element Output Bounds for the Incompressible Navier-Stokes Equations: Application to a Natural Convection Problem.” *J. Comput. Phys.*, 172:401–425, 2001.
- [92] Maday, Y., Kaber, S. M. O., and Tadmor, E. “Legendre Pseudospectral Viscosity Method for Nonlinear Conservation Laws.” *SIAM Journal on Numerical Analysis*, 30(2):321–342, 1993.
- [93] Maday, Y. and Tadmor, E. “Analysis of the Spectral Vanishing Viscosity Method for Periodic Conservation Laws.” *SIAM Journal on Numerical Analysis*, 26(4):854–870, 1989.
- [94] Mavriplis, C. “Adaptive mesh strategies for the spectral element method.” *Comput. Methods Appl. Mech. Engrg.*, 116:77–86, 1994.
- [95] Nastase, C. R. and Mavriplis, D. J. “High-order discontinuous Galerkin methods using an *hp*-multigrid approach.” *Journal of Computational Physics*, 213(1):330–357, 2006.
- [96] Nemec, M., Aftosmis, M. J., and Wintzer, M. “Adjoint-Based Adaptive Mesh Refinement for Complex Geometries.” AIAA Paper 2008-725, 2008.
- [97] Nielsen, E. *FUN2D/3D Fully Unstructured Navier-Stokes User Manual*. NASA Langley Research Center, Computational Modeling and Simulation Branch, Virginia. <http://fun3d.larc.nasa.gov>.
- [98] Nompelis, I., Drayna, T. W., and Candler, G. V. “Development of a Hybrid Unstructured Implicit Solver for the Simulation of Reacting Flows over Complex Geometries.” AIAA Paper 2004-2227, 2004.
- [99] Oberai, A. A. and Wanderer, J. “A Dynamic Approach for Evaluating Parameters in a Numerical Method.” *International Journal for Numerical Methods in Engineering*, 62:50–71, 2005.

- [100] Oberai, A. A. and Wanderer, J. “A Dynamic Multiscale Viscosity method for the Spectral Approximation of Conservation Laws.” *Computer Methods in Applied Mechanics and Engineering*, 195:1778–1792, 2006.
- [101] Oliver, T. A. “Multigrid Solution for High-Order Discontinuous Galerkin Discretizations of the Compressible Navier-Stokes Equations.” Master’s thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, August 2004.
- [102] Oliver, T. A. *A High-Order, Adaptive, Discontinuous Galerkin Finite Element Method for the Reynolds-Averaged Navier-Stokes Equations*. PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2008.
- [103] Oliver, T. A. and Darmofal, D. L. “An analysis of dual consistency for discontinuous Galerkin discretization of source terms.” ACDL report, 2007.
- [104] Ozcer, I. A. “Sonic Boom Prediction Using Euler / Full Potential Methodology.” AIAA Paper 2007-369, 2007.
- [105] Page, J. A. and Plotkin, K. J. “An Efficient Method for Incorporating Computational Fluid Dynamics Into Sonic Boom Prediction.” AIAA Paper 91-3275, 1991.
- [106] Peraire, J., Vahdati, M., Morgan, K., and Zienkiewicz, O. C. “Adaptive Remeshing for Compressible Flow Computations.” *J. Comput. Phys.*, 72:449–466, 1987.
- [107] Persson, P.-O. and Peraire, J. “An Efficient Low Memory Implicit DG Algorithm for Time Dependent Problems.” AIAA Paper 2006-0113, 2006.
- [108] Persson, P.-O. and Peraire, J. “Sub-Cell Shock Capturing for Discontinuous Galerkin Methods.” AIAA Paper 2006-112, 2006.
- [109] Pierce, N. A. and Giles, M. B. “Adjoint and defect error bounding and correction for functional estimates.” *Journal of Computational Physics*, 200(2):769–794, 2004.
- [110] Pilon, A. R. “Spectrally Accurate Prediction of Sonic Boom Signals.” *AIAA Journal*, 45: 2149–2156, 2007.
- [111] Pittman, J., Bartolotta, P. A., and Mansour, N. N. “Fundamental Aeronautics Hypersonics Project Reference Document.” NASA, 2006. <http://www.aeronautics.nasa.gov/fap/documents.html>.
- [112] Plotkin, K. J. “State of the art of sonic boom modeling.” *Journal of the Acoustical Society of America*, 111:530–536, 2002.
- [113] Qiu, J. and Shu, C.-W. “Hermite WENO schemes and their application as limiters for Runge-Kutta discontinuous Galerkin method: One-dimensional case.” *Journal of Computational Physics*, 193(1):115–135, 2004.
- [114] Qiu, J. and Shu, C.-W. “Hermite WENO schemes and their application as limiters for Runge-Kutta discontinuous Galerkin method II: Two dimensional case.” *Computers and Fluids*, 34: 642–663, 2005.
- [115] Quattrocchi, D. J. *Hypersonic heat transfer and anisotropic visualization with a higher-order discontinuous Galerkin finite element method*. Masters thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 2006.
- [116] Rallabhandi, S. K. and Mavris, D. N. “A New Approach for Incorporating Computational Fluid Dynamics into Sonic Boom Prediction.” AIAA Paper 2006-3312, 2006.
- [117] Rasmussen, M. *Hypersonic Flow*. John Wiley & Sons, Inc., New York, 1994.

- [118] Reed, W. H. and Hill, T. R. "Triangular mesh methods for the neutron transport equation." Report Technical Report LA-UR-73-479, Los Alamos Scientific Laboratory, 1973.
- [119] Remaki, M. and Habashi, W. G. "A discontinuous Galerkin method/HLLC solver for the Euler equations." *International Journal for Numerical Methods in Fluids*, 43(12):1391–1405, 2003.
- [120] Roe, P. L. "Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes." *J. Comput. Phys.*, 43:357–372, 1981.
- [121] Rudy, D. H., Thomas, J. L., Kumar, A., Gnoffo, P. A., and Chakravarthy, S. R. "Computation of laminar hypersonic compression-corner flows." *AIAA Journal*, 29(7):1108–1113, 1991.
- [122] Schwanenberg, D. and Harms, M. "Discontinuous Galerkin finite-element method for transcritical two-dimensional shallow water flows." *Journal of Hydraulic Engineering*, 130(5):412–421, 2004.
- [123] Seebass, R. and Argrow, B. "Sonic Boom Minimization Revisited." AIAA Paper 98-2956, 1998.
- [124] Sherwin, S. J. and Karniadakis, G. E. "A new triangular and tetrahedral basis for high-order finite element methods." *Int. J. Num. Meth. Eng.*, 38:3775–3802, 1995.
- [125] Shu, C. W. and Osher, S. "Efficient implementation of essentially non-oscillatory shock-capturing schemes." *J. Comput. Phys.*, 77:439–471, 1988.
- [126] Si, H. *TetGen, A Quality Tetrahedral Mesh Generator and Three-Dimensional Delaunay Triangulator*. Weierstrass Institute for Applied Analysis and Stochastic, Berlin, Germany, 2006. <http://tetgen.berlios.de>.
- [127] Süli, E., Houston, P., and Senior, B. "*hp*-Discontinuous Galerkin finite element methods for hyperbolic problems: error analysis and adaptivity." In Baines, M. J., editor, *Numerical Methods for Fluid Dynamics VII*. ICFD, Oxford, 2001.
- [128] Szabo, B. A. "Estimation and Control of Error Based on p Convergence." In Babuška, I., Zienkiewicz, O. C., Gago, J., and de A. Oliveira, E. R., editors, *Accuracy Estimates and Adaptive Refinements in Finite Element Computations*, pages 61–78. John Wiley & Sons Ltd., 1986.
- [129] Tadmor, E. "Convergence of Spectral Methods for Nonlinear Conservation Laws." *SIAM Journal on Numerical Analysis*, 26(1):30–44, 1989.
- [130] Tadmor, E. "Shock Capturing by the Spectral Viscosity Method." *Computer Methods in Applied Mechanics and Engineering*, 80:197–208, 1990.
- [131] Tadmor, E. "Total Variation and Error Estimates for Spectral Viscosity Approximations." *Mathematics of Computation*, 60:245–256, 1993.
- [132] van Leer, B. "Towards the ultimate conservative difference scheme. II - Monotonicity and conservation combined in a second order scheme." *Journal of Computational Physics*, 14:361–370, 1974.
- [133] van Leer, B. "Towards the ultimate conservative difference scheme. III - Upstream-centered finite-difference schemes for ideal compressible flow." *Journal of Computational Physics*, 23: 263–275, 1977.
- [134] van Leer, B. "Towards the ultimate conservative difference scheme. IV - A new approach to numerical convection." *Journal of Computational Physics*, 23:276–299, 1977.

- [135] van Leer, B. "Towards the ultimate conservative difference scheme. V - A second-order sequel to Godunov's method (for ideal compressible flow)." *Journal of Computational Physics*, 32: 101–136, 1979.
- [136] Venditti, D. A. and Darmofal, D. L. "Grid Adaptation for Functional Outputs: Application to Two-Dimensional Inviscid Flows." *J. Comput. Phys.*, 176:40–69, 2002.
- [137] Venditti, D. A. and Darmofal, D. L. "Anisotropic Grid Adaptation for Functional Outputs: Application to Two-Dimensional Viscous Flows." *J. Comput. Phys.*, 187:22–46, 2003.
- [138] Venkatakrishnan, V., Allmaras, S. R., Kamenetskii, D. S., and Johnson, F. T. "Higher order schemes for the compressible Navier-Stokes equations." AIAA Paper 2003-3987, 2003.
- [139] von Neumann, J. and Richtmyer, R. "A method for the numerical calculation of hydrodynamic shocks." *Journal of Applied Physics*, 21:232–237, 1950.
- [140] Wang, L. and Mavriplis, D. J. "Implicit Solution of the Unsteady Euler Equations for High-Order Accurate Discontinuous Galerkin Discretizations." AIAA Paper 2006-0109, 2006.
- [141] Warbuton, T. C., Lomtev, I., Du, Y., Sherwin, S. J., and Karniadakis, G. E. "Galerkin and Discontinuous Galerkin spectral/hp Methods." *Comput. Methods Appl. Mech. Engrg.*, 175: 343–359, 1999.
- [142] Warren, G. P., Anderson, W. K., Thomas, J. L., and Krist, S. L. "Grid Convergence for Adaptive Methods." AIAA Paper 91-1592, 1991.
- [143] Wilcoxson, M. and Manousiouthakis, V. "On an Implicit ENO Scheme." *Journal of Computational Physics*, 115:376–389, 1994.
- [144] Yang, J.-Y., Perng, Y.-C., and Yen, R.-H. "Implicit Weighted Essentially Nonoscillatory Schemes for the Compressible Navier-Stokes Equations." *AIAA Journal*, 39(11):2082–2090, 2001.
- [145] Yoon, S., Gnoffo, P. A., White, J. A., and Thomas, J. L. "Computational Challenges in Hypersonic Flow Simulations." AIAA Paper 2007-4265, 2007.
- [146] Zienkiewicz, O. C. and Zhu, J. Z. "Adaptivity and Mesh Generation." *International Journal for Numerical Methods in Engineering*, 32:783–810, 1991.